Advanced Journey With

Gustavo A. Hoffmann Robert A. Duff

A Flight in Progress





Advanced Journey With Ada: A Flight In Progress

Release 2025-06

Gustavo A. Hoffmann and Robert A. Duff

Jun 27, 2025

CONTENTS

1 Types 5 1.1 Names 5 1.1.1 Direct names 5 1.1.2 Other kinds of names 6 1.2 Objects 9 1.2.1 Constant and variable objects 11 1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4.1 Enumerations as functions 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.2 Root Types 38 1.8 Type view 38 1.8 Type view 38 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 58 1.9.4 Conversion of	I	Dat	ta types 3	}
1.1.1 Direct names 5 1.2.2 Other kinds of names 6 1.2 Otherts 9 1.2.1 Constant and variable objects 11 1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4.1 Enumerations 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.3 Non-Record Private Types 31 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplet types 36 1.8 INon-Record Private Types 38 <	1	Туре	25	5
1.1.2 Other kinds of names 6 1.2 Objects 9 1.2.1 Constant and variable objects 11 1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations as functions 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 38 1.9.1 View conversion 45 <t< th=""><th></th><th>1.1</th><th>Names</th><th>5</th></t<>		1.1	Names	5
1.2 Objects 9 1.2.1 Constant and variable objects 11 1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations as functions 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 32 1.8 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversions 54 1.9.3 Implicit conversions 54			1.1.1 Direct names	5
1.2.1 Constant and variable objects 11 1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4.1 Enumerations 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 INon-Record Private Types 41 1.9 Type view 38 1.8.1 Non-Record Private Types 58 <t< td=""><td></td><td></td><td>1.1.2 Other kinds of names</td><td>5</td></t<>			1.1.2 Other kinds of names	5
1.2.1 Constant and variable objects 11 1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4.1 Enumerations 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 INon-Record Private Types 41 1.9 Type view 38 1.8.1 Non-Record Private Types 58 <t< td=""><td></td><td>1.2</td><td>Objects</td><td>)</td></t<>		1.2	Objects)
1.2.2 View of an object 12 1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8.1 Non-Record Private Types 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Outlified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User				
1.2.3 Named numbers 13 1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations 20 1.4.1 Enumeration as functions 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values <td></td> <td></td> <td>,</td> <td></td>			,	
1.3 Scalar Types 14 1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal and Root Types 28 1.5.2 Root Types 28 1.5.3 Non-Record Private Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 58 1.9.4 Conversion 44 1.10.1 Verifying subtypes 64				
1.3.1 Ranges 15 1.3.2 Predecessor and Successor 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.3 Root Types 28 1.5.4 Non-Record Private Types 28 1.5.7 Incomplete types 31 1.7 Incomplete types 31 1.7 Incomplete types 31 1.7 Incomplete types 31 1.8 INon-Record Private Types 41 1.9 Type view 38 1.8.1 Non-Record Private Types 59 1.10 Qualified Expressions 54 1.9		13		
1.3.2 Predecessor and Successor. 15 1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations as functions 20 1.4.1 Enumeration overloading 20 1.4.2 Enumeration overloading 20 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.3 Non-Record Private Types 28 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 31 1.7 Incomplete types 31 1.8 Non-Record Private Types 41 1.9 Type view 38 1.8.1 Non-Record Private Types 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 59 1.10 Qualified Expressions 64 1.10.1 Verifying subtypes		1.5		
1.3.3 Scalar To String Conversion 18 1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations as functions 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 31 1.9 Type conversion 45 1.9.1 View conversion 45 1.9.2 View conversion 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64			5	
1.3.4 Width attribute 19 1.3.5 Other attributes 20 1.4 Enumerations as functions 20 1.4.1 Enumeration overloading 20 1.4.2 Enumeration overloading 20 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation				
1.3.5 Other attributes 20 1.4 Enumerations 20 1.4.1 Enumeration overloading 23 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 28 1.5.2 Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 31 1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversions 54 1.9.2 View conversions 54 1.9.3 Implicit conversions 54 1.9.3 Implic			5	
1.4 Enumerations 20 1.4.1 Enumeration overloading 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplet types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81				
1.4.1 Enumeration s as functions 20 1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 31 1.8 Type view 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Deferred Constants 65 1.12 Deferred Constants 68 1.13 User-defined literals 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80		1 /		
1.4.2 Enumeration overloading 23 1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3.1 Sizes 81 2.3.2 Component size 81		1.4		
1.4.3 Position and Internal Code 26 1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 65 1.13 User-defined literals 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3.1 Sizes 81 2.3.2 Component size 81 <td></td> <td></td> <td></td> <td></td>				
1.5 Universal and Root Types 28 1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 31 1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 45 1.9.2 View conversions 58 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.10.1 Verifying subtypes 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 <td< td=""><td></td><td></td><td></td><td></td></td<>				
1.5.1 Universal Types 28 1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 65 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3.1 Size attribute and aspect 81 2.3.2 Component size 81		1 5		
1.5.2 Root Types 29 1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83		1.5		
1.6 Definite and Indefinite Subtypes 31 1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
1.7 Incomplete types 36 1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 54 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 65 1.13 User-defined literals 70 2 Types and Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 81		1.0		
1.8 Type view 38 1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 45 1.9.3 Implicit conversions 54 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.10.1 Verifying subtypes 64 1.11 Default initial values 65 1.12 Deferred Constants 65 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
1.8.1 Non-Record Private Types 41 1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 54 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
1.9 Type conversion 45 1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83		1.8		
1.9.1 Value conversion 45 1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 81				
1.9.2 View conversion 54 1.9.3 Implicit conversions 58 1.9.4 Conversion of other types 59 1.10 Qualified Expressions 64 1.10.1 Verifying subtypes 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83		1.9		
1.9.3 Implicit conversions581.9.4 Conversion of other types591.10 Qualified Expressions641.10.1 Verifying subtypes641.11 Default initial values651.12 Deferred Constants681.13 User-defined literals702 Types and Representation792.1 Enumeration Representation Clauses792.2 Data Representation802.3 Sizes812.3.1 Size attribute and aspect812.3.2 Component size83				
1.9.4 Conversion of other types591.10 Qualified Expressions641.10.1 Verifying subtypes641.11 Default initial values651.12 Deferred Constants681.13 User-defined literals702 Types and Representation792.1 Enumeration Representation Clauses792.2 Data Representation802.3 Sizes812.3.1 Size attribute and aspect812.3.2 Component size83				
1.10 Qualified Expressions641.10.1 Verifying subtypes641.11 Default initial values651.12 Deferred Constants681.13 User-defined literals702 Types and Representation792.1 Enumeration Representation Clauses792.2 Data Representation802.3 Sizes812.3.1 Size attribute and aspect812.3.2 Component size83				
1.10.1 Verifying subtypes 64 1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
1.11 Default initial values 65 1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 70 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83		1.10		
1.12 Deferred Constants 68 1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
1.13 User-defined literals 70 2 Types and Representation 79 2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
2 Types and Representation792.1 Enumeration Representation Clauses792.2 Data Representation802.3 Sizes812.3.1 Size attribute and aspect812.3.2 Component size83				3
2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83		1.13	User-defined literals)
2.1 Enumeration Representation Clauses 79 2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83	2	Type	es and Representation 79)
2.2 Data Representation 80 2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
2.3 Sizes 81 2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
2.3.1 Size attribute and aspect 81 2.3.2 Component size 83				
2.3.2 Component size		2.5		

	2.4 2.5 2.6	Alignment89Overlapping Storage92Packed Representation94
		2.6.1 Trade-offs
	2.7	Record Representation and storage clauses
		2.7.1 Storage Place Attributes
		2.7.2 Using Representation Clauses
		2.7.3 Derived Types And Representation Clauses
		2.7.4 Representation on Bit Level
	2.8	Changing Data Representation
		2.8.1 Restrictions
		Valid Attribute
		Unchecked Union
	2.11	Addresses
		2.11.1 Address attribute
		2.11.2 Address aspect
		2.11.3 Address comparison
		2.11.4 Address to integer conversion
		2.11.5 Address arithmetic
	2.12	Discarding names
3	Shar	red variable control 135
5		Volatile
	3.2	Independent
	3.3	Atomic
	3.4	Full-access only
	5.4	3.4.1 Nonatomic full-access
		3.4.2 Atomic full-access
		3.4.3 Comparison: full-access and non-full-access types
	3.5	Atomic operations
	0.0	3.5.1 Atomic Exchange
4	Reco	
	4.1	Default Initialization
		4.1.1 Dependencies
		4.1.2 Initialization Order
		4.1.2 Initialization Order 172 4.1.3 Evaluation 173
		4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174
		4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176
	4.2	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177
	4.2 4.3	4.1.2 Initialization Order1724.1.3 Evaluation1734.1.4 Defaults and object declaration1744.1.5 Advanced Usages176Mutually dependent types177Null records179
		4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176 Mutually dependent types 177 Null records 179 4.3.1 Simple Prototyping 180
		4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176 Mutually dependent types 177 Null records 179 4.3.1 Simple Prototyping 180 4.3.2 Extending the prototype 182
		4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176 Mutually dependent types 177 Null records 179 4.3.1 Simple Prototyping 180 4.3.2 Extending the prototype 182 4.3.3 More complex applications 184
		4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API184
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186
		4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants188
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants1884.4.1Known and unknown discriminant parts188
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants1884.4.1Known and unknown discriminant parts1884.4.2Discriminant as constant property189
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants1884.4.1Known and unknown discriminant parts1884.4.2Discriminant as constant property1894.4.3Private types190
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants1884.4.1Known and unknown discriminant parts1884.4.2Discriminant as constant property1894.4.4Object declaration191
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants1884.4.1Known and unknown discriminant parts1884.4.2Discriminant as constant property1894.4.3Private types1904.4.4Object declaration1914.4.5Object assignments193
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records1884.4.1Known and unknown discriminant parts1884.4.2Discriminant as constant property1894.4.3Private types1904.4.4Object declaration1914.4.5Object assignments1934.4.6Discriminant type194
	4.3	4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176 Mutually dependent types 177 Null records 179 4.3.1 Simple Prototyping 180 4.3.2 Extending the prototype 182 4.3.3 More complex applications 184 4.3.4 Implementing the API 184 4.3.5 Tagged null records 188 4.4.1 Known and unknown discriminant parts 188 4.4.2 Discriminant as constant property 189 4.4.3 Private types 190 4.4.4 Object declaration 191 4.4.5 Object assignments 193 4.4.6 Discriminant type 194 4.4.7 Default values 197
	4.3	4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176 Mutually dependent types 177 Null records 179 4.3.1 Simple Prototyping 180 4.3.2 Extending the prototype 182 4.3.3 More complex applications 184 4.3.4 Implementing the API 184 4.3.5 Tagged null records 188 4.4.1 Known and unknown discriminant parts 188 4.4.2 Discriminant as constant property 189 4.4.3 Private types 190 4.4.4 Object declaration 191 4.4.5 Object assignments 193 4.4.6 Discriminant type 194 4.4.7 Default values 197 4.4.8 Derived types and subtypes 201
	4.3	4.1.2Initialization Order1724.1.3Evaluation1734.1.4Defaults and object declaration1744.1.5Advanced Usages176Mutually dependent types177Null records1794.3.1Simple Prototyping1804.3.2Extending the prototype1824.3.3More complex applications1844.3.4Implementing the API1844.3.5Tagged null records186Record discriminants1884.4.1Known and unknown discriminant parts1884.4.2Discriminant as constant property1894.4.3Private types1904.4.4Object declaration1914.4.5Object assignments1934.4.6Discriminant type1944.4.7Default values1974.4.8Derived types and subtypes201Discriminant constraints and operations211
	4.3	4.1.2 Initialization Order 172 4.1.3 Evaluation 173 4.1.4 Defaults and object declaration 174 4.1.5 Advanced Usages 176 Mutually dependent types 177 Null records 179 4.3.1 Simple Prototyping 180 4.3.2 Extending the prototype 182 4.3.3 More complex applications 184 4.3.4 Implementing the API 184 4.3.5 Tagged null records 188 4.4.1 Known and unknown discriminant parts 188 4.4.2 Discriminant as constant property 189 4.4.3 Private types 190 4.4.4 Object declaration 191 4.4.5 Object assignments 193 4.4.6 Discriminant type 194 4.4.7 Default values 197 4.4.8 Derived types and subtypes 201

	4.6	Unknown discriminants	219
	4.7	4.6.3 Derived types	225
	4.8	Variant parts	230
		4.8.3 Ensuring valid information	234 236
	4.9	Per-Object Expressions	243
5	Agg	regates 2	247
		Container Aggregates	249
		5.2.1 <>	255
	5.3 5.4	Full coverage rules for Aggregates	260 262
		5.4.1 Positional and named array aggregates5.4.2 Null array aggregate5.4.3 , <>, others	265 266
		5.4.4	268 269
	5.5	5.4.8 <> and default values 5.4.8 <> and default values 5.5.1 Assignments to objects of derived types	276 281
		5.5.2 Example: Points	281 283 284
	FC	5.5.5 with others5.5.6 with null record5.5.7 Extension aggregates and descendent types	286 286
	5.6	Delta Aggregates	287 290
		5.6.3 Delta Aggregates for Arrays	
6	Arra 6.1	Array constraints	295
	6.2	6.1.3 Constrained array types Multidimensional Arrays 6.2.1 Unconstrained Multidimensional Arrays	298 303
	6.3	 6.2.2 Arrays of arrays Derived array types and array subtypes 6.3.1 Derived array types 6.3.2 Array subtypes 	307 307
7	Stri		300 811
-		Character and String Literals	311 311

	7.2	Wide and Wide-Wide Strings	314
		7.2.1 Text I/O	
		7.2.2 Wide and Wide-Wide String Handling	
		7.2.3 Bounded and Unbounded Wide and Wide-Wide Strings	
	7.3	String Encoding	
		7.3.1 UTF-8 encoding and decoding	
		7.3.2 UTF-8 size and length	
	7 4	7.3.3 UTF-16 encoding and decoding	
	7.4	UTF-8 applications	
		7.4.1 UTF-8 encoding in source-code files7.4.2 Parsing UTF-8 files for Wide-Wide-String processing	
	7.5	Image attribute	
	1.5	7.5.1 Overview	
		7.5.2 Type'Image and Obj'Image	
		7.5.3 Wider versions of Image	
		7.5.4 Image attribute for non-scalar types	
		7.5.5 Image attribute for tagged types	
		7.5.6 Image attribute for task and protected types	
	7.6	Put Image aspect	
		7.6.1 Overview	
		7.6.2 Complete Example of Put Image	
		7.6.3 Relation to the Image attribute	
		7.6.4 Put_Image and derived types	348
		7.6.5 Put_Image and tagged types	
	7.7	Universal text buffer	
		7.7.1 Overview	
		7.7.2 Additional procedures	352
8	Nun	nerics	355
8		nerics Numeric Literals	355 355
8	Nun 8.1	Numeric Literals	355
8		Numeric Literals	355 355
8		Numeric Literals8.1.1 Classification8.1.2 Features and Flexibility	355 355 357
8	8.1	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types	355 355 357 362
8	8.1	Numeric Literals8.1.1 Classification8.1.2 Features and Flexibility	355 355 357 362 362
8	8.1	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer	355 355 357 362 362 362 371
8	8.1 8.2	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal Fixed	355 355 357 362 362 371 376
8	8.1 8.2	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types	355 355 357 362 362 371 376 377 381
8	8.18.28.3	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute	355 355 357 362 371 376 377 381 381
8	8.18.28.3	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute	355 355 357 362 371 376 377 381 381 382
8	8.18.28.38.4	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types	355 355 357 362 371 376 377 381 381 382 384
8	8.18.28.3	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular typesAttributes of Floating-Point Types	355 355 357 362 371 376 371 376 377 381 381 382 384 386
8	8.18.28.38.4	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular typesAttributes of Floating-Point Types8.5.1 Representation-oriented attributes	355 355 357 362 371 376 377 381 381 381 382 384 386 386
8	8.18.28.38.48.5	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular typesAttributes of Floating-Point Types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes	355 355 357 362 371 376 377 381 381 381 382 384 386 386 386 391
8	8.18.28.38.4	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.5.2 Primitive function attributesAttributes of Fixed-Point types	355 355 357 362 371 376 377 381 381 381 381 382 384 386 386 391 399
8	8.18.28.38.48.5	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.5.1 Attributes of Fixed-Point types8.5.1 Attributes of Fixed-Point types8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.1 Attributes of Fixed-Point types8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.3 Primitive function attributes8.5.4 Determine function attributes8.5.5 Primitive function attributes8.5.1 Attributes of ordinary and decimal fixed-point types	355 355 357 362 371 376 377 381 381 381 382 384 386 386 391 399 399
8	 8.1 8.2 8.3 8.4 8.5 8.6 	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.5.2 Primitive function attributes8.5.1 Attributes of Floating-Point Types8.5.2 Primitive function attributes8.5.2 Primitive function attributes8.5.2 Attributes of Fixed-Point types8.6.1 Attributes of ordinary and decimal fixed-point types8.6.2 Attributes of decimal fixed-point types	355 355 357 362 371 376 377 381 381 381 382 384 386 386 391 399 399 406
8	8.18.28.38.48.5	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.5.2 Primitive function attributes8.6.1 Attributes of ordinary and decimal fixed-point types8.6.2 Attributes of decimal fixed-point types8.6.2 Attributes of decimal fixed-point typesBig Numbers	355 355 357 362 371 376 377 381 381 381 382 384 386 386 391 399 399 406 408
8	 8.1 8.2 8.3 8.4 8.5 8.6 	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.5.2 Primitive function attributes8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.3 Attributes of Fixed-Point types8.5.4 Attributes of Fixed-Point types8.5.5 Attributes of Fixed-Point types8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.3 Representation-oriented attributes8.5.4 Attributes of Fixed-Point types8.5.5 Representation-oriented attributes8.5.7 Primitive function attributes8.5.8 Attributes of Fixed-Point types8.5.1 Attributes of ordinary and decimal fixed-point types8.5.2 Attributes of ordinary and decimal fixed-point types8.5.3 Attributes of ordinary and decimal fixed-point types8.5.4 Attributes of decimal fixed-point types8.5.5 Attributes of decimal fixed-point types8.5.7 Attributes of decimal fixed-point types8.5.8 Attributes of decimal fixed-point types8.5.9 Attributes of decimal fixed-point types8.5.1 Overview	355 355 357 362 371 376 377 381 381 381 382 384 386 386 386 391 399 399 406 408 409
8	 8.1 8.2 8.3 8.4 8.5 8.6 	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal Fixed8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.5.2 Attributes of Floating-Point Types8.5.1 Attributes of Fixed-Point types8.5.2 Primitive function attributes8.5.3 Mumbers8.5.4 Attributes of ordinary and decimal fixed-point types8.5.7 Attributes of Jeacimal fixed-point types8.5.1 Attributes of ordinary and decimal fixed-point types8.5.2 Primitive function attributes8.5.3 Representation-oriented attributes8.5.4 Attributes of ordinary and decimal fixed-point types8.5.7 Attributes of Jeacimal fixed-point types	355 355 357 362 371 376 377 381 381 381 381 382 384 386 386 391 399 399 406 408 409 410
8	 8.1 8.2 8.3 8.4 8.5 8.6 	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.6.1 Attributes of ordinary and decimal fixed-point types8.6.2 Attributes of decimal fixed-point types8.7.1 Overview8.7.2 Factorial8.7.3 Conversions	355 355 357 362 371 376 377 381 381 381 382 384 386 386 391 399 399 406 408 409 410 413
8	 8.1 8.2 8.3 8.4 8.5 8.6 	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.6.1 Attributes of ordinary and decimal fixed-point types8.6.2 Attributes of decimal fixed-point types8.6.1 Attributes of decimal fixed-point types8.7.1 Overview8.7.2 Factorial8.7.4 Other features of big integers	355 355 357 362 371 376 377 381 381 382 384 386 386 391 399 399 406 408 409 410 413 420
8	 8.1 8.2 8.3 8.4 8.5 8.6 	Numeric Literals8.1.1 Classification8.1.2 Features and FlexibilityUniversal Numeric Types8.2.1 Universal Real and Integer8.2.2 Universal FixedBase types8.3.1 BaseAttributes of Modular Types8.4.1 Modulus Attribute8.4.2 Mod Attribute8.4.3 Operations on modular types8.5.1 Representation-oriented attributes8.5.2 Primitive function attributes8.6.1 Attributes of ordinary and decimal fixed-point types8.6.1 Attributes of decimal fixed-point types8.6.2 Attributes of decimal fixed-point types8.7.1 Overview8.7.2 Factorial8.7.4 Other features of big integers	355 355 357 362 371 376 377 381 381 382 384 386 386 391 399 399 406 408 409 410 413 420 421

II Control Flow

_	_		
9	Ехрі		27
	9.1	Expressions: Definition	427
		9.1.1 Relations and simple expressions	427
		9.1.2 Numeric expressions	430
		9.1.3 Other expressions	
		9.1.4 Parenthesized expression	
	0.2		
	9.2	Conditional Expressions	
	9.3	Quantified Expressions	
	9.4	Declare Expressions	
		9.4.1 Restrictions in the declarative part	442
	9.5	Reduction Expressions	
	0.0	9.5.1 Value sequences	
		9.5.2 Custom reducers	
		9.5.3 Other accumulator types	448
10			51
	10.1	Simple and Compound Statements	451
	10.2	Labels	451
		10.2.1 Labels and goto statements	
		10.2.2 Use-case: Continue	
		10.2.3 Labels and compound statements	
		Exit loop statement	
	10.4	If, case and loop statements	457
		10.4.1 Case statements and expressions	458
	10.5	Block Statements	
		Extended return statement	
	10.0	10.6.1 Other usages of extended return statements	
			+05
	Curk		6 F
11			65
	11.1	Parameter Modes and Associations	
		11.1.1 Formal Parameter Modes	
		11.1.2 By-copy and by-reference	465
		11.1.3 Bounded errors	470
		11.1.4 Aliased parameters	
		11.1.5 Parameter Associations	
	11 2		
	11.2	Operators	
		11.2.1 User-defined operators	
		Expression functions	
	11.4	Overloading	487
	11.5	Operator Overloading	491
		Operator Overriding	
		Nonreturning procedures	
		Inline subprograms	
	11.9	Null Procedures	
		11.9.1 Null procedures and overriding	501
12			05
	12.1	Classification of Errors	505
		12.1.1 Compilation errors	505
		12.1.2 Runtime errors	
		12.1.3 Bounded errors	
		12.1.4 Erroneous execution	
		Asserts	
	12.3	Assertion policies	510
		Checks and exceptions	
		12.4.1 Access Check	
		12.4.2 Discriminant Check	
			111

	12.4.3 Division Check	517
	12.4.4 Index Check	517
	12.4.5 Length Check	518
		519
	12.4.7 Range Check	
	12.4.8 Tag Check	
	12.4.9 Accessibility Check	
	12.4.10Allocation Check	
	12.4.1 Elaboration Check	
	12.4.12Program_Error_Check	
	12.4.13Storage Check	
	12.4.14Tasking_Check	529
12.5	Ada.Exceptions package	530
	12.5.1 Retrieving exception information	531
	12.5.2 Collecting exceptions	
	12.5.3 Debugging exceptions in the GNAT toolchain	
	Exception renaming	
	Out and Uninitialized	
	Suppressing checks	
	12.8.1 pragma Suppress	
	12.8.2 pragma Unsuppress	545

III Modular programming

547

13 Packages	549
13.1 Package renaming	549
13.1.1 Grouping packages	549
13.1.2 Child of renamed package	551
13.1.3 Backwards-compatibility via renaming	551
13.2 Private packages	
13.2.1 Declaration and usage	553
13.2.2 Private sibling packages	
13.2.3 Outside the package tree	557
13.3 Private with clauses	560
13.3.1 Definition and usage	
13.3.2 Referring to private child package	
13.4 Limited Visibility	
13.4.1 Limited visibility and private with clauses	
13.4.2 Limited visibility and other elements	
13.5 Visibility	
13.5.1 Automatic visibility	
13.5.2 With clauses and visibility	
13.5.3 Circular dependency	
13.5.4 Private packages	
13.6 Use type clause	
13.6.1 Another use clause example	
13.6.2 Visibility and Readability	
13.6.3 use type	
13.6.4 use all type	
13.7 Use clauses and naming conflicts	
13.7.1 Code example	
13.7.2 Naming conflict	
13.7.3 Circumventing naming conflicts	580
14Subprograms and Modularity	585
14.1 Private subprograms	
14.1.1 Private subprograms of a package	
14.1.2 Private subprograms and private packages	

IV Resource Management

591

15Access Types 593	3
15.1 Access types: Terminology 59)3
15.1.1 Access type, designated subtype and profile	3
15.1.2 Access object and designated object	
15.1.3 Access value and designated value	4
15.2 Access types: Allocation	95
15.2.1 Pool-specific access types)7
15.2.2 Multiple allocation	
15.3 Discriminants as Access Values 60	
15.3.1 Unconstrained type as designated subtype 60	
15.3.2 Whole object assignments	
15.4 Parameters as Access Values	
15.4.1 Changing the referenced object	
15.4.2 Replace the access value	
15.4.3 Side-effects on designated objects	.4
15.5 Self-reference	
15.6 Mutually dependent types using access types	3
15.7 Dereferencing	
15.7.1 Implicit Dereferencing	5
15.8 Ragged arrays	
15.8.1 Uniform multidimensional arrays	0
15.8.2 Non-uniform multidimensional array	
15.9 Aliasing	
15.9.1 Aliased objects	
15.9.2 Aliased components	
15.9.3 Aliased parameters	.3
15.10Accessibility Levels and Rules: An Introduction	
15.10. Lifetime of objects	
15.10.2Accessibility Levels	
15.10.3Accessibility Rules	
15.10.4Accessibility rules on parameters	
15.10. Dangling References	
15.11Unchecked Access	
15.12 Jnchecked Deallocation	
15.12. Unchecked Deallocation and Dangling References	
15.12. Dereferencing dangling references	
15.12. Restrictions for Ada. Unchecked_Deallocation	
15.13Null & Not Null Access	
15.14Design strategies for access types	
15.14. Abstract data type for access types	
15.14.2Controlled type for access types	
15.15Access to subprograms	
15.15.15tatic vs. dynamic calls	
15.15.2Access to subprogram declaration	
15.15.30bjects of access-to-subprogram type	
15.15.4Components of access-to-subprogram type	
15.15.5Access-to-subprogram as discriminant types	
15.15.75electing subprograms	
15.15.8Null exclusion	
15.16Accessibility Rules and Access-To-Subprograms	
15.16.1Unchecked Access	
15.17Access and Address	
15.17.14ddress and access conversion	
	-

16 Ano	nymous Access Types	711
16.1	Named and Anonymous Access Types	. 711
	16.1.1 Relation to named types	. 712
	16.1.2 Benefits of anonymous access types	
16.2	Anonymous Access-To-Óbject Types	
	16.2.1 Not Null Anonymous Access-To-Object Types	
	16.2.2 Drawbacks of Anonymous Access-To-Object Types	
163	Access discriminants	
10.5	16.3.1 Default Value of Access Discriminants	. 725
	16.3.2 Benefits of Access Discriminants	
1.0.4	16.3.3 Preventing dangling pointers	
	Self-reference	
	Mutually dependent types using anonymous access types	
16.6	Access parameters	
	16.6.1 Interfacing To Other Languages	
	16.6.2 Inherited Primitive Operations For Tagged Types	
16.7	' User-Defined References	
	16.7.1 Dereferencing of tagged types	. 746
	16.7.2 Simple container	. 747
16.8	Anonymous Access Types and Accessibility Rules	. 753
	16.8.1 Conversions between Anonymous and Named Access Types	
	16.8.2 Accessibility rules on access parameters	
16.9	Anonymous Access-To-Subprograms	
	16.9.1 Examples of anonymous access-to-subprogram usage	
	16.9.2 Application of anonymous access-to-subprogram types	
	16.9.3 Readability	
16 1	.OAccessibility Rules and Anonymous Access-To-Subprograms	
10.1	16.10. Named vs. anonymous access-to-subprograms	
	16.10.2Named vs. anonymous access-to-subprograms as parameters	
	16.10.3terator	
	10.10.3(e)a(o)	. //2
17 Lim	ited Types	779
	Assignment and equality	
	17.1.1 Assignments	782
	17.1.2 Equality	
17 2	Limited private types	
17.2	17.2.1 Non-Record Limited Types	
	17.2.2 Partial and full view of limited types	
	17.2.3 Limited and nonlimited in full view	
	17.2.4 Limited private component	
170	17.2.5 Tagged limited private types	
	Explicitly limited types	
	Subtypes of Limited Types	
17.5	Deriving from limited types	
	17.5.1 Deriving from limited private types	
	17.5.2 Deriving from non-explicitly limited private types	
17.6	immutably Limited Types	
	17.6.1 Non immutably limited types	. 807
17.7	' Limited Types with Discriminants	. 808
	17.7.1 Default Expressions	. 809
	17.7.2 Limited private type with unknown discriminants	. 812
17.8	Record components of limited type	
) Limited types and aggregates	
	17.9.1 Full coverage rules for limited types	
171	.Constructor functions for limited types	
	.Return objects	
1/.1	17.11. Extended return statements for limited types	
	17.11.2nitialization and function return	
		. UZJ

17.12Building objects from constructors	
18 Controlled Types 83	-
18.1 Overview	5
18.1.1 Lifetime of objects	5
18.1.2 Initialization of objects	6
18.1.3 Controlled objects	
18.1.4 Adjustment of controlled objects	
18.1.5 Limited controlled types	4
18.1.6 Simple Example with ID	
18.2 Initialization	
18.2.1 Subcomponents	
18.2.2 Components with access discriminants	
18.2.3 Task activation	
18.3 Assignment	
18.3.1 Assignment using anonymous object	
18.3.2 Adjustment of subcomponents	
18.4 Finalization	
18.4.1 Normal and abnormal completion	
18.4.2 Finalization via unchecked deallocation	
18.4.3 Subcomponents	
18.4.4 Components with access discriminants	
18.5 Controlled Types and Exception Handling	
18.5.1 Exception raising in Initialize	
18.5.2 Bounded errors of controlled types	
18.5.3 Memory allocation and exceptions	
18.6 Applications of Controlled Types	
18.6.1 Encapsulating access type handling	
18.6.2 Encapsulating file handling	

Copyright © 2019 - 2024, AdaCore

This book is published under a CC BY-SA license, which means that you can copy, redistribute, remix, transform, and build upon the content for any purpose, even commercially, as long as you give appropriate credit, provide a link to the license, and indicate if changes were made. If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original. You can find license details on this page¹



🛕 Warning

This is work in progress!

Information in this document is subject to change at any time without prior notification.

1 Note

The code examples in this course use a 50-column limit, which greatly improves the readability of the code on devices with a small screen size. This constraint, however, leads to an unusual coding style. For instance, instead of calling Put_Line in a single line, we have this:

```
Put_Line
  (" is in the northeast quadrant");
```

or this:

Note that typical Ada code uses a limit of at least 79 columns. Therefore, please don't take the coding style from this course as a reference!

\rm 1 Note

Each code example from this book has an associated "code block metadata", which contains the name of the "project" and an MD5 hash value. This information is used to identify a single code example.

You can find all code examples in a zip file, which you can download from the learn website². The directory structure in the zip file is based on the code block metadata. For example, if you're searching for a code example with this metadata:

- Project: Courses.Intro_To_Ada.Imperative_Language.Greet
- MD5: cba89a34b87c9dfa71533d982d05e6ab

you will find it in this directory:

¹ http://creativecommons.org/licenses/by-sa/4.0

projects/Courses/Intro_To_Ada/Imperative_Language/Greet/ cba89a34b87c9dfa71533d982d05e6ab/

In order to use this code example, just follow these steps:

- 1. Unpack the zip file;
- 2. Go to target directory;
- 3. Start GNAT Studio on this directory;
- 4. Build (or compile) the project;
- 5. Run the application (if a main procedure is available in the project).

This course will teach you advanced topics of the Ada programming language. The Introduction to Ada³ course is a prerequisite for this course.

This document was written by Gustavo A. Hoffmann, with major contributions from Robert A. Duff. The document also includes contributions from Franco Gasperoni, Gary Dismukes, Patrick Rogers, and Robert Dewar.

These contributions are clearly indicated in the document, together with the original publication source.

Special thanks to Patrick Rogers for all comments and suggestions. In particular, thanks for sharing the training slides on access types: many ideas from those slides were integrated into this course.

This document was reviewed by Patrick Rogers and Tucker Taft.

CHANGELOG

Changes are being tracked on the CHANGELOG page.

² https://learn.adacore.com/zip/learning-ada_code.zip

³ https://learn.adacore.com/courses/intro-to-ada/index.html#intro-ada-course-index

Part I Data types

TYPES

1.1 Names

In simple terms, a "name" can be an identifier, i.e. the *name* that we use to refer to an object or a subprogram, for example. This is what we call a *direct name*. However, in Ada, a name can also refer to other language constructs, as we discuss later on in this section.

In the Ada Reference Manual

• 4.1 Name⁴

1.1.1 Direct names

Direct names are the simplest form of names in Ada. They can be either identifiers or operator symbols.

Identifiers

An identifier — as the term implies — is a (direct) name that we use to *identify* an object, a subprogram, a type, and so on. When specifying an identifier, we aren't limited to ASCII⁵ characters: we can use a subset of the Unicode⁶ standard.

1 For further reading...

To be more precise, the Normalization Form KC of the Unicode standard is applied to identifiers. You can find more information about it in the Unicode Standard Annex $#15^7$.

For example:

Listing 1: show	_identifiers.adb
-----------------	------------------

```
procedure Show_Identifiers is
1
               ~~~~
                     ~~~~~~
   - -
2
                  identifier
   - -
3
4
      type New Integer is new
5
6
              identifier
7
         Integer;
8
```

- ⁴ http://www.ada-auth.org/standards/22rm/html/RM-4-1.html
- ⁵ https://en.wikipedia.org/wiki/ASCII
- ⁶ https://en.wikipedia.org/wiki/Universal_Coded_Character_Set
- ⁷ https://unicode.org/reports/tr15/

```
~~~~
9
            identifier
10
       - -
11
      Something_Important : New_Integer;
12
13
       - -
           identifier
       - -
14
                               ~~~~~~
       - -
15
                                identifier
16
   begin
17
     null;
18
  end Show_Identifiers;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Names.Identifiers
MD5: e427d3e5fe5f549df593b5e5941cf2ba
```

In this example, we see the following identifiers: Show_Identifiers (subprogram), New_Integer (type), Integer (type), and Something_Important (object).

Operator symbols

The set of operator symbols that we can use is restricted to the following symbols or reserved words specified in the Ada language:

Operator kind	Operators
Logical operators	and, or, xor
Relational operators	=, /=, <, <= >, >=
Binary adding operators	+, -, &
Unary adding operators	+, -
multiplying opertors	*,/, mod , rem
Highest precedence operators	**, abs, not

In the Ada Reference Manual

4.5 Operators and Expression Evaluation⁸

1.1.2 Other kinds of names

In addition to direct names, we have the following kinds of names: *explicit dereferences* (page 623), indexed components, slices, selected components, attribute references, *type conversions* (page 45), function calls, character literals, *qualified expressions* (page 64), *generalized references* (page 744), and target name.

Let's see an example of some of them:

Listing 2: show_other_names.adb

```
pragma Ada_2022;
procedure Show_Other_Names is
type Integer_Access is
```

⁸ http://www.ada-auth.org/standards/22rm/html/RM-4-5.html

```
7
       type Integer_Array is array
8
         (Positive range <>) of Integer;
9
10
       type New_Integer is new
11
         Integer;
12
13
       function Zero
14
         return New_Integer is
15
           (0);
16
17
       subtype Sub_Integer is
18
19
         Integer;
20
       type Rec is record
21
          Val : Integer := 0;
22
       end record;
23
24
       type ABC_Enum is
25
         ('A', 'B', 'C');
26
27
       IA : Integer_Access := new Integer;
28
       Arr : Integer_Array (1 .. 5) :=
29
30
                (others => 0);
         : Rec;
       R
31
       NI : New_Integer;
32
       SI : Sub_Integer;
33
       Е
          : ABC_Enum := 'A';
34
   begin
35
       R.Val := IA.all;
36
                ~~~~~
       - -
37
       -- explicit dereference
38
39
       R.Val := Arr (1);
40
                 ~~~~~
41
       - -
       - -
             indexed component
42
43
       Arr (1 .. 2) := Arr (3 .. 4);
44
                        ~~~~~~~~~
       - -
45
       - -
                             slice
46
47
       Arr (1 .. 2) := (others => R.Val);
48
                                     ~~~~
49
       - -
       - -
                        selected component
50
51
       R.Val := Arr'Size;
52
                 ~~~~~
53
       - -
       -- attribute reference
54
55
       NI := New_Integer (IA.all);
56
             _____
57
       - -
       - -
               type conversion
58
59
       NI := Zero;
60
              ~~/
61
       - -
       -- function call
62
63
       E := 'A';
64
             ~~~
65
       - -
       -- character literal
66
```

access Integer;

6

```
67
       IA.all := Sub_Integer (R.Val);
68
69
       - -
       - -
                  qualified expression
70
71
       R.Val := @ + 1;
72
73
       - -
          target name
       - -
74
       - -
75
       -- equivalent to:
76
               R.Val := R.Val + 1;
77
78
   end Show_Other_Names;
79
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Names.Other_Names
MD5: 8063a4c9ff7a01ff7a69454fae096089
```

In this example, we see instances of the following kinds of names:

- explicit dereference: IA.all;
- indexed components: Arr (1);
- slices: Arr (1 .. 2), Arr (3 .. 4);
- selected components: R.Val;
- attribute references: Arr'Size;
- type conversions: New_Integer (IA.all);
- function calls: Zero;
- character literals: 'A';
- qualified expressions: Sub_Integer (R.Val);
- target name: @.

1 In the Ada Reference Manual

- 4.1 Name⁹
- 4.1.1 Indexed Components¹⁰
- 4.1.2 Slices¹¹
- 4.1.3 Selected Components¹²
- 4.1.4 Attributes¹³
- 4.1.5 User-Defined References¹⁴
- 4.6 Type Conversions¹⁵
- 4.7 Qualified Expressions¹⁶
- 5.2.1 Target Name Symbols¹⁷

1.2 Objects

The term *object* may be misleading for readers that have a strong background in objectoriented programming. Moreover, its meaning can vary depending on the context. Therefore, it's important to define what we mean by *objects* when focusing on Ada programming.

In computer science, the term $object^{18}$ can refer to a piece of data stored in memory — but it can also refer to a table or a form in a database. Also, even when we define the term object as data in memory, we can still classify programming languages as object-based¹⁹ or object-oriented²⁰ languages.

1 Important

In object-oriented programming, an object belongs to a *class* of objects. In Ada, objects of this kind are called tagged objects. Note, however, that we can have objects that don't belong to a class of objects: those are called *untagged* objects.

In the context of Ada programming, an object is an "entity that contains a value, and is either a constant or a variable" — according to the Ada Reference Manual. In other words, any constants or variables that we declare in Ada source code are objects. In addition, there are other examples of objects that don't originate from object declarations:

1	<pre>procedure Show_Objects is</pre>
2	
3	type New_Integer is new
4	Integer;
5	
6	type Integer_Array is
7	array (<mark>Positive</mark> range <>) of
8	Integer;
9	
10	procedure Dummy (Obj : Integer)
11	is null;
12	
13	object
14	
15	task type TT is
16	<pre>entry Start (Id : Integer); ^^</pre>
17	
18	object
19	end TT;
20	tool hada TT da
21	task body TT is
22	begin
23	accept Start (Id : Integer) do
24	null;

Listing 3: show objects.adb

(continues on next page)

⁹ http://www.ada-auth.org/standards/22rm/html/RM-4-1.html

¹⁰ http://www.ada-auth.org/standards/22rm/html/RM-4-1-1.html

¹¹ http://www.ada-auth.org/standards/22rm/html/RM-4-1-2.html

¹² http://www.ada-auth.org/standards/22rm/html/RM-4-1-3.html

¹³ http://www.ada-auth.org/standards/22rm/html/RM-4-1-4.html ¹⁴ http://www.ada-auth.org/standards/22rm/html/RM-4-1-5.html

¹⁵ http://www.ada-auth.org/standards/22rm/html/RM-4-6.html

¹⁶ http://www.ada-auth.org/standards/22rm/html/RM-4-7.html

¹⁷ http://www.ada-auth.org/standards/22rm/html/RM-5-2-1.html

¹⁸ https://en.wikipedia.org/wiki/Object (computer science)

¹⁹ https://en.wikipedia.org/wiki/Object-based language

²⁰ https://en.wikipedia.org/wiki/Object-oriented programming

```
end Start;
25
       end TT;
26
27
       function Add_One (V : Integer)
28
29
       - -
             view of an object
30
       - -
                          return Integer is
31
       begin
32
          return V + 1;
33
                  ~~~~
          - -
34
          - -
                  object
35
       end Add_One;
36
37
       Arr : Integer_Array (1 .. 10);
38
           ~~~~/
39
       - -
       -- object
40
41
       NI : New_Integer;
42
   begin
43
       Arr (1 .. 3) := (others => 1);
44
           ~~~~~
45
       - -
       -- object
46
                         ^^^^
       - -
47
       - -
                            object
48
49
       NI := New_Integer (Arr (1));
50
              ~~~~~
51
       - -
              object
52
       - -
53
       for I in Arr'Range loop
54
55
       -- object
56
57
          Arr (I) := Add_One (Arr (I));
58
59
          - -
                                object
          - -
60
       end loop;
61
   end Show_Objects;
62
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Objects.Object_Examples
MD5: edf9eab70ec0ecce90ef71591324ac94

As we can see in this code example a formal parameter of a subprogram or an entry is also an object — in addition, so are *value conversions* (page 45), the result returned by a function, the result of evaluating an *aggregate* (page 247), loop parameters, *arrays* (page 295), or the slices of arrays objects, or the components of composite objects.

Other examples of objects include:

- the object created via a view conversion (page 54);
- a dereference (page 623) of an access-to-variable (page 637) value;
- · the return object of a function;
- a choice parameter of an exception handler²¹.

 $^{^{21}\} https://learn.adacore.com/courses/intro-to-ada/chapters/exceptions.html \# intro-ada-handling-an-exception$

In the Ada Reference Manual

3.3 Objects and Named Numbers²²

1.2.1 Constant and variable objects

Objects can be classified as constant and variable objects. When declaring objects, the distinction is clear:

Listing 4: show_objects.adb

```
1 procedure Show_Objects is
2 Const : constant Integer := 42;
3 Var : Integer := 0;
4 begin
5 null;
6 end Show_Objects;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Objects.Object_Declaration_Examples
MD5: 16b4d9546e9c05443ced05c7f6608cc9
```

In this example, Const is a constant object, while Var is a variable object.

In addition to this, constant objects include:

- the discriminant component (page 188) of a variable discriminant;
- a formal parameter or generic formal object of mode in.

On the other hand, variable objects include:

- the object created via a view conversion (page 54) of a variable;
- a dereference (page 623) of an access-to-variable (page 637) value.

For example:

Listing 5: show_objects.adb

```
procedure Show_Objects is
1
2
      type Device (Id : Positive) is
3
       record
4
          Value : Integer;
5
      end record;
6
7
      type Device_Access is
8
         access all Device;
9
10
      Dev : aliased Device (99);
11
12
       -- Discriminant `Id` is a
13
          constant object.
       - -
14
15
       -- `Dev` is a variable object,
16
       -- though.
17
18
      Dev_Acc : Device_Access := Dev'Access;
19
20
```

(continues on next page)

²² http://www.ada-auth.org/standards/22rm/html/RM-3-3.html

```
procedure Process (D : Device) is
21
         null;
22
23
       - -
                  constant object
       - -
24
   begin
25
       Dev.Value := 0;
26
27
       - -
       -- variable object
28
29
       Dev_Acc.all.Value := 1;
30
31
       - -
           variable object
32
33
   end Show_Objects;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Objects.Object_Examples
MD5: c0d1386a37e5ed31f0d3163fadbb1b30
```

In this example, we see that Dev is a variable object, while its Id discriminant is a constant object. In addition, the Dev_Acc.all dereference is a variable object. Finally, the in parameter of procedure Process is a constant object.

```
1 In the Ada Reference Manual
```

- 3.3 Objects and Named Numbers²³
- 3.3.1 Object Declarations²⁴

1.2.2 View of an object

As we've just seen, an object can be either constant or variable. In addition, the *view* of an object is classified as constant or variable as well.

Before we start, note that the classification of an object as constant or variable doesn't directly imply how its view is classified. You may, for example, expect that a constant object has a constant view, but this is not necessarily the case, as we discuss in this section. (In fact, a constant object only has a constant view if it doesn't have a part that has a variable view.)

A part of an object has a variable view if it is of *immutably limited type* (page 805), *controlled type* (page 838), *private type* (page 38), or private extension. In that sense, if any of those parts with variable view exist in a constant object, then we say that the *whole object* has a variable view. Only if a constant object doesn't have *any* parts with variable view, then this object has a constant view.

In contrast, variable objects always have a variable view.

Let's see an example:

Listing 6: devices.ads

```
1 package Devices is
2
3 type Device_Settings is
4 record
5 Started : Boolean;
```

(continues on next page)

²³ http://www.ada-auth.org/standards/22rm/html/RM-3-3.html
 ²⁴ http://www.ada-auth.org/standards/22rm/html/RM-3-3-1.html

```
end record;
6
7
       type Device (Id : Positive) is
8
         private;
9
10
       function Init (Id : Positive)
11
                        return Device;
12
13
   private
14
15
       type Device (Id : Positive) is
16
         null record;
17
18
       function Init (Id : Positive)
19
                        return Device is
20
          (Device'(Id => Id));
21
22
   end Devices;
23
```

Listing 7: show object view.adb

```
with Devices; use Devices;
1
2
   procedure Show_Object_View is
З
      Dev
             : constant Device := Init (5);
4
      -- Constant object with
5
      -- variable view.
6
7
      Default : constant Device_Settings
8
                   := (Started => False);
9
      -- Constant object with
10
      -- constant view.
11
12
      Settings : Device_Settings;
13
14
   begin
15
      Settings := (Started => True);
16
   end Show Object View;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Objects.Object_View
MD5: b9a56ee937e71c728bac116f21d98742
```

In this example, both Default_S and Dev are constant objects. However, they have different views: while Default_S has a constant view because it doesn't have any parts with variable view, Dev has a variable view because it's a private type. Finally, as expected, Settings has a variable view because it's a variable object.

1.2.3 Named numbers

In addition to objects, we can have named numbers. Those aren't objects, but rather *names* (page 5) that we assign to numeric values. For example:

Listing 8: show_named_number.adb

```
procedure Show_Named_Number is
Pi : constant := 3.1415926535;
```

- 5 begin
- 6 null;
- 7 end Show_Named_Number;

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Objects.Named_Number MD5: ee6808bb7ecb7fef687831f53a8b6668

In this example, Pi is a named number.

A named number is always known at compilation time. Also, it doesn't have a type associated with it. In fact, its type is called universal real or universal integer — depending on the number being a real or integer number. (In this specific case, Pi is a universal real number.) We talk about *universal types* (page 28) later on in this chapter and about *universal real and integer types* (page 362) in another chapter.

6	In	the	Ada	Reference Manua	
---	----	-----	-----	------------------------	--

3.3.2 Number Declarations²⁵

1.3 Scalar Types

In general terms, scalar types are the most basic types that we can get. As we know, we can classify them as follows:

Category	Discrete	Numeric
Enumeration	Yes	No
Integer	Yes	Yes
Real	No	Yes

Many attributes exist for scalar types. For example, we can use the Image and Value attributes to convert between a given type and a string type. The following table presents the main attributes for scalar types:

Category	At- tribute	Returned value
Ranges	First	First value of the discrete subtype's range.
	Last	Last value of the discrete subtype's range.
	Range	Range of the discrete subtype (corresponds to Subtype'First Subtype'Last).
Iterators	Pred	Predecessor of the input value.
	Succ	Successor of the input value.
Comparison	Min	Minimum of two values.
	Max	Maximum of two values.
String con-	Image	String representation of the input value.
version	Value	Value of a subtype based on input string.

We already discussed some of these attributes in the Introduction to Ada course (in the

²⁵ http://www.ada-auth.org/standards/22rm/html/RM-3-3-2.html

sections about range and related attributes²⁶ and image attribute²⁷). In this section, we'll discuss some aspects that have been left out of the previous course.

1 In the Ada Reference Manual

```
    3.5 Scalar types<sup>28</sup>
```

1.3.1 Ranges

We've seen that the First and Last attributes can be used with discrete types. Those attributes are also available for real types. Here's an example using the **Float** type and a subtype of it:

```
Listing 9: show first last real.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
    procedure Show First Last Real is
3
        subtype Norm is Float range 0.0 .. 1.0;
4
    begin
5
        Put_Line ("Float'First: " & Float'First'Image);
Put_Line ("Float'Last: " & Float'Last'Image);
6
7
        Put_Line ("Norm'First: " & Norm'First'Image);
Put_Line ("Norm'Last: " & Norm'Last'Image);
8
9
    end Show First Last Real;
10
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Scalar_Types.Ranges_Real_Types
MD5: 89745a94fbdc41a2880ba14e50401acb

Runtime output

Float'First:	-3.40282E+38
Float'Last:	3.40282E+38
Norm'First:	0.00000E+00
Norm'Last:	1.00000E+00

This program displays the first and last values of both the **Float** type and the Norm subtype. In the case of the **Float** type, we see the full range, while for the Norm subtype, we get the values we used in the declaration of the subtype (i.e. 0.0 and 1.0).

1.3.2 Predecessor and Successor

We can use the Pred and Succ attributes to get the predecessor and successor of a specific value. For discrete types, this is simply the next discrete value. For example, Pred (2) is 1 and Succ (2) is 3. Let's look at a complete source-code example:

Listing 10:	show succ	pred di	screte.adb

```
with Ada.Text_I0; use Ada.Text_I0;
```

```
1
2
3
```

```
procedure Show_Succ_Pred_Discrete is
   type State is (Idle, Started,
```

²⁶ https://learn.adacore.com/courses/intro-to-ada/chapters/arrays.html#intro-ada-range-attribute
²⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html#

intro-ada-image-attribute

²⁸ http://www.ada-auth.org/standards/22rm/html/RM-3-5.html

```
Processing, Stopped);
5
6
      Machine_State : constant State := Started;
7
8
      I : constant Integer := 2;
9
10
   begin
      Put_Line ("State
                                              1.1
11
                 & Machine_State'Image);
12
      Put_Line ("State'Pred (Machine_State): "
13
                 & State'Pred (Machine_State)'Image);
14
      Put_Line ("State'Succ (Machine_State):
15
                 & State'Succ (Machine_State)'Image);
16
      Put_Line ("-----");
17
18
                                   : "
      Put_Line ("I
19
                 & I'Image);
20
      Put_Line ("Integer'Pred (I): "
21
                 & Integer'Pred (I)'Image);
22
      Put_Line ("Integer'Succ (I): "
23
                 & Integer'Succ (I)'Image);
24
   end Show_Succ_Pred_Discrete;
25
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Scalar_Types.Show_Succ_Pred_Discrete
MD5: e11d0f50105864fdc1594b3bb72d927e

Runtime output

State : STARTED
State'Pred (Machine_State): IDLE
State'Succ (Machine_State): PROCESSING
.....
I : 2
Integer'Pred (I): 1
Integer'Succ (I): 3

In this example, we use the Pred and Succ attributes for a variable of enumeration type (State) and a variable of **Integer** type.

We can also use the Pred and Succ attributes with real types. In this case, however, the value we get depends on the actual type we're using:

- for fixed-point types, the value is calculated using the smallest value (Small), which is derived from the declaration of the fixed-point type;
- for floating-point types, the value used in the calculation depends on representation constraints of the actual target machine.

Let's look at this example with a decimal type (Decimal) and a floating-point type (My_Float):

Listing 11: show_succ_pred_real.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Show_Succ_Pred_Real is
subtype My_Float is
Float range 0.0 .. 0.5;
type Decimal is
delta 0.1 digits 2
```

```
range 0.0 .. 0.5;
9
10
      D : Decimal;
11
      N : My_Float;
12
13
   begin
      Put_Line ("---- DECIMAL -----");
14
      Put_Line ("Small: " & Decimal'Small'Image);
15
      Put_Line ("----- Succ ------");
16
      D := Decimal'First;
17
      loop
18
          Put Line (D'Image);
19
          D := Decimal'Succ (D);
20
21
          exit when D = Decimal'Last;
22
       end loop;
23
      Put_Line ("----- Pred ------");
24
25
      D := Decimal'Last;
26
      loop
27
          Put_Line (D'Image);
28
          D := Decimal'Pred (D);
29
30
          exit when D = Decimal'First;
31
      end loop;
32
33
      Put_Line ("========");
34
      Put_Line ("---- MY_FLOAT ----");
35
      Put_Line ("----- Succ ------");
36
      N := My_Float'First;
37
       for I in 1 .. 5 loop
38
          Put_Line (N'Image);
39
          N := My_Float'Succ (N);
40
       end loop;
41
      Put_Line ("----- Pred ------");
42
43
      for I in 1 .. 5 loop
44
          Put_Line (N'Image);
45
          N := My_Float'Pred (N);
46
       end loop;
47
   end Show_Succ_Pred_Real;
48
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Scalar_Types.Show_Succ_Pred_Real
MD5: f426d6539c3ce863101f1e6afb21c08f

Runtime output

0.1
MY_FLOAT 0.000000E+00 1.40130E-45 2.80260E-45 4.20390E-45
4.20390E-43 5.60519E-45 7.00649E-45 5.60519E-45 4.20390E-45 2.80260E-45 1.40130E-45

=

As the output of the program indicates, the smallest value (see Decimal'Small in the example) is used to calculate the previous and next values of Decimal type.

In the case of the My_Float type, the difference between the current and the previous or next values is 1.40130E-45 (or 2^{-149}) on a standard PC.

1.3.3 Scalar To String Conversion

We've seen that we can use the Image and Value attributes to perform conversions between values of a given subtype and a string:

Listing 12:	show	image	value	attr.adb
-------------	------	-------	-------	----------

```
with Ada.Text_IO; use Ada.Text_IO;
procedure Show_Image_Value_Attr is
I : constant Integer := Integer'Value ("42");
begin
Put_Line (I'Image);
end Show_Image_Value_Attr;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Scalar_Types.Image_Value_Attr MD5: 9daa13b1f05511fac7e108eb9b8eefa7

Runtime output

42

The Image and Value attributes are used for the **String** type specifically. In addition to them, there are also attributes for different string types — namely **Wide_String** and Wide_Wide_String. This is the complete list of available attributes:

Conversion type	Attribute	String type
Conversion to string	Image	String
	Wide_Image	Wide_String
	Wide_Wide_Image	Wide_Wide_String
Conversion to subtype	Value	String
	Wide_Value	Wide_String
	Wide_Wide_Value	Wide_Wide_String

We discuss more about Wide_String and Wide_Wide_String in another section (page 314).

1.3.4 Width attribute

When converting a value to a string by using the Image attribute, we get a string with variable width. We can assess the maximum width of that string for a specific subtype by using the Width attribute. For example, **Integer**'Width gives us the maximum width returned by the Image attribute when converting a value of **Integer** type to a string of **String** type.

This attribute is useful when we're using bounded strings in our code to store the string returned by the Image attribute. For example:

Listing 13: show_width_attr.adb

```
with Ada.Text I0;
                                use Ada.Text I0;
1
   with Ada.Strings;
                                use Ada.Strings;
2
   with Ada.Strings.Bounded;
3
4
   procedure Show Width Attr is
5
      package B Str is new
6
         Ada.Strings.Bounded.Generic Bounded Length
7
           (Max => Integer'Width);
8
      use B_Str;
9
10
      Str_I : Bounded_String;
11
12
      I : constant Integer := 42;
13
      J : constant Integer := 103;
14
   begin
15
      Str_I := To_Bounded_String (I'Image);
16
      Put_Line ("Value:
17
                 & To_String (Str_I));
18
      Put_Line ("String Length:
19
                 & Length (Str I) 'Image);
20
      Put_Line ("----");
21
22
      Str_I := To_Bounded_String (J'Image);
23
      Put_Line ("Value:
24
                 & To_String (Str_I));
25
      Put_Line ("String Length:
26
                 & Length (Str_I)'Image);
27
   end Show Width Attr;
28
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Scalar_Types.Width_Attr MD5: 82cff0cf4fecfdecce3020135cf98fd2

Runtime output

Value: 42 String Length: 3 ----Value: 103 String Length: 4

In this example, we're storing the string returned by Image in the Str_I variable of Bounded_String type.

Similar to the Image and Value attributes, the Width attribute is also available for string types other than **String**. In fact, we can use:

- the Wide_Width attribute for strings returned by Wide_Image; and
- the Wide_Wide_Width attribute for strings returned by Wide_Wide_Image.

1.3.5 Other attributes

The Base attribute is specific for numeric types. We discuss this topic *in another chapter* (page 377).

1.4 Enumerations

We've introduced enumerations back in the Introduction to Ada course²⁹. In this section, we'll discuss a few useful features of enumerations, such as enumeration renaming, enumeration overloading and representation clauses.

1 In the Ada Reference Manual

• 3.5.1 Enumeration Types³⁰

1.4.1 Enumerations as functions

If you have used programming language such as C in the past, you're familiar with the concept of enumerations being constants with integer values. In Ada, however, enumerations are not integers. In fact, they're actually parameterless functions! Let's consider this example:

```
Listing 14: days.ads
```

```
package Days is
1
2
       type Day is (Mon, Tue, Wed,
3
                     Thu, Fri,
4
                     Sat, Sun);
5
6
       -- Essentially, we're declaring
7
          these functions:
       - -
8
9
       - -
           function Mon return Day;
10
       - -
       -- function Tue return Day;
11
           function Wed return Day;
       - -
12
           function Thu return Day;
13
       - -
           function Fri return Day;
14
       - -
           function Sat return Day;
       - -
15
          function Sun return Day;
       - -
16
17
   end Days;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_As_Function
MD5: fa3e58b58edffa5a3e04b060a7f8cb8b

In the package Days, we're declaring the enumeration type Day. When we do this, we're essentially declaring seven parameterless functions, one for each enumeration. For example, the Mon enumeration corresponds to **function** Mon return Day. You can see all seven function declarations in the comments of the example above.

Note that this has no direct relation to how an Ada compiler generates machine code for

²⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/strongly_typed_language.html# intro-ada-enum-types

³⁰ http://www.ada-auth.org/standards/22rm/html/RM-3-5-1.html

enumeration. Even though enumerations are parameterless functions, a typical Ada compiler doesn't generate function calls for code that deals with enumerations.

Enumeration renaming

The idea that enumerations are parameterless functions can be used when we want to rename enumerations. For example, we could rename the enumerations of the Day type like this:

```
Listing 15: enumeration_example.ads
```

```
package Enumeration Example is
1
2
      type Day is (Mon, Tue, Wed,
3
                    Thu, Fri,
4
                    Sat, Sun);
5
6
      function Monday
                         return Day renames Mon;
7
      function Tuesday return Day renames Tue;
8
      function Wednesday return Day renames Wed;
9
      function Thursday return Day renames Thu;
10
      function Friday return Day renames Fri;
11
      function Saturday return Day renames Sat;
12
      function Sunday
                         return Day renames Sun;
13
14
  end Enumeration Example;
15
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Renaming
MD5: e2e12bb3bfcb0b6e94769ced9a4b80f9

Now, we can use both Monday or Mon to refer to Monday of the Day type:

Listing 16: show_renaming.adb

```
use Ada.Text I0;
   with Ada.Text IO;
1
   with Enumeration_Example; use Enumeration_Example;
2
3
   procedure Show Renaming is
4
      D1 : constant Day := Mon;
5
      D2 : constant Day := Monday;
6
   begin
7
      if D1 = D2 then
8
         Put Line ("D1 = D2");
9
          Put_Line (Day'Image (D1)
10
                    & " =
11
                    & Day'Image (D2));
12
      end if:
13
   end Show Renaming;
14
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Renaming MD5: 2d7177def2c9e9fb11c7dc5e036c3be3

Runtime output

D1 = D2MON = MON

When running this application, we can confirm that D1 is equal to D2. Also, even though

we've assigned Monday to D2 (instead of Mon), the application displays Mon = Mon, since Monday is just another name to refer to the actual enumeration (Mon).

```
1 Hint
   If you just want to have a single (renamed) enumeration visible in your application -
   and make the original enumeration invisible —, you can use a separate package. For
   example:
                           Listing 17: enumeration_example.ads
   package Enumeration Example is
1
2
      type Day is (Mon, Tue, Wed,
3
                    Thu, Fri,
4
                    Sat, Sun);
5
6
   end Enumeration Example;
7
                           Listing 18: enumeration renaming.ads
   with Enumeration Example;
1
2
   package Enumeration Renaming is
3
4
      subtype Day is Enumeration_Example.Day;
5
6
                          return Day renames
7
      function Monday
        Enumeration_Example.Mon;
8
      function Tuesday
                         return Day renames
9
        Enumeration_Example.Tue;
10
      function Wednesday return Day renames
11
        Enumeration_Example.Wed;
12
      function Thursday return Day renames
13
        Enumeration Example.Thu;
14
      function Friday
                         return Day renames
15
        Enumeration Example.Fri;
16
      function Saturday return Day renames
17
        Enumeration_Example.Sat;
18
                         return Day renames
      function Sunday
19
        Enumeration_Example.Sun;
20
21
   end Enumeration Renaming;
22
                              Listing 19: show renaming.adb
   with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Enumeration_Renaming;
3
   use Enumeration_Renaming;
4
5
   procedure Show Renaming is
6
      D1 : constant Day := Monday;
7
   begin
8
      Put Line (Day'Image (D1));
9
   end Show Renaming;
10
   Code block metadata
   Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Renaming
   MD5: 87fe75026f0fc118921eaee45fe55a8a
   Runtime output
```

MON

Note that the call to Put_Line still display Mon instead of Monday.

1.4.2 Enumeration overloading

Enumerations can be overloaded. In simple terms, this means that the same name can be used to declare an enumeration of different types. A typical example is the declaration of colors:

Listing 20: colors.ads

```
package Colors is
1
2
       type Color is
3
         (Salmon,
4
5
          Firebrick,
          Red,
6
          Darkred,
7
          Lime,
8
          Forestgreen,
9
          Green,
10
          Darkgreen,
11
          Blue,
12
          Mediumblue,
13
          Darkblue);
14
15
       type Primary Color is
16
         (Red,
17
          Green,
18
          Blue);
19
20
   end Colors;
21
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Overloading
MD5: b808f90d9164f044b6b7a8931863726f

Note that we have Red as an enumeration of type Color and of type Primary_Color. The same applies to Green and Blue. Because Ada is a strongly-typed language, in most cases, the enumeration that we're referring to is clear from the context. For example:

Listing 21: red_colors.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   with Colors;
                      use Colors;
2
3
   procedure Red Colors is
4
      C1 : constant Color
                                    := Red;
5
      -- Using Red from Color
6
7
      C2 : constant Primary_Color := Red;
8
      -- Using Red from Primary Color
9
   begin
10
      if C1 = Red then
11
         Put_Line ("C1 = Red");
12
      end if;
13
      if C2 = Red then
14
         Put_Line ("C2 = Red");
15
```

16 end if;

17 end Red_Colors;

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Overloading
MD5: dd590eab88164773e974e748d77a51af

Runtime output

C1 = RedC2 = Red

When assigning Red to C1 and C2, it is clear that, in the first case, we're referring to Red of Color type, while in the second case, we're referring to Red of the Primary_Color type. The same logic applies to comparisons such as the one in **if** C1 = Red: because the type of C1 is defined (Color), it's clear that the Red enumeration is the one of Color type.

Enumeration subtypes

Note that enumeration overloading is not the same as enumeration subtypes. For example, we could define the following subtype:

Listing 22: colors-shades.ads

```
1 package Colors.Shades is
2
3 subtype Blue_Shades is
4 Colors range Blue .. Darkblue;
5
6 end Colors.Shades;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Overloading
MD5: 9c13508bda487cae02dbf8b403271540

In this case, Blue of Blue_Shades and Blue of Colors are the same enumeration.

Enumeration ambiguities

A situation where enumeration overloading might lead to ambiguities is when we use them in ranges. For example:

Listing 23: colors.ads

1	package Colors is			
2				
3	type <mark>Colo</mark> r is			
4	(Salmon,			
5	Firebrick,			
6	Red,			
7	Darkred,			
8	Lime,			
9	Forestgreen,			
10	Green,			
11	Darkgreen,			
12	Blue,			
13	Mediumblue,			

```
14 Darkblue);
15
16 type Primary_Color is
17 (Red,
18 Green,
19 Blue);
20
21 end Colors;
```

Listing 24: color loop.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Colors:
                      use Colors:
2
З
   procedure Color Loop is
4
   begin
5
      for C in Red .. Blue loop
6
7
          ERROR: range is ambiguous!
8
         Put Line (Color'Image (C));
9
      end loop;
10
   end Color_Loop;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Ambiguities
MD5: 82d0d3f28f1faf6b296a4f44db71f41b

Build output

```
color_loop.adb:6:17: error: ambiguous bounds in range of iteration
color_loop.adb:6:17: error: possible interpretations:
color_loop.adb:6:17: error: type "Primary_Color" defined at colors.ads:16
color_loop.adb:6:17: error: type "Color" defined at colors.ads:3
color_loop.adb:6:17: error: ambiguous bounds in discrete range
color_loop.adb:9:30: error: expected type "Color" defined at colors.ads:3
color_loop.adb:9:30: error: found type "Primary_Color" defined at colors.ads:16
gprbuild: *** compilation phase failed
```

Here, it's not clear whether the range in the loop is of Color type or of Primary_Color type. Therefore, we get a compilation error for this code example. The next line in the code example — the one with the call to Put_Line — gives us a hint about the developer's intention to refer to the Color type. In this case, we can use qualification — for example, Color' (Red) — to resolve the ambiguity:

Listing 25: color_loop.adb

```
with Ada.Text_IO; use Ada.Text IO;
1
  with Colors;
                     use Colors;
2
3
   procedure Color_Loop is
4
  begin
5
      for C in Color'(Red) .. Color'(Blue) loop
6
         Put_Line (Color'Image (C));
7
      end loop;
8
  end Color Loop;
9
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Ambiguities
MD5: c3e946d330bb6aed258bcd005a540794

Runtime output

RED DARKRED LIME FORESTGREEN GREEN DARKGREEN BLUE

Note that, in the case of ranges, we can also rewrite the loop by using a range declaration:

Listing 26: color_loop.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
  with Colors;
                  use Colors;
2
3
  procedure Color Loop is
4
  begin
5
      for C in Color range Red .. Blue loop
6
         Put Line (Color'Image (C));
7
      end loop;
8
  end Color_Loop;
9
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Ambiguities
MD5: 23f8db4fcb5710f7bda6b511234e0448

Runtime output

RED DARKRED LIME FORESTGREEN GREEN DARKGREEN BLUE

Alternatively, Color **range** Red .. Blue could be used in a subtype declaration, so we could rewrite the example above using a subtype (such as Red_To_Blue) in the loop:

Listing 27: color loop.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
  with Colors;
                     use Colors;
2
3
   procedure Color_Loop is
4
      subtype Red_To_Blue is Color range Red .. Blue;
5
6
   begin
      for C in Red_To_Blue loop
7
         Put_Line (Color'Image (C));
8
      end loop;
9
  end Color_Loop;
10
```

1.4.3 Position and Internal Code

As we've said above, a typical Ada compiler doesn't generate function calls for code that deals with enumerations. On the contrary, each enumeration has values associated with it, and the compiler uses those values instead.

Each enumeration has:

- a position value, which is a natural value indicating the position of the enumeration in the enumeration type; and
- an internal code, which, by default, in most cases, is the same as the position value.

Also, by default, the value of the first position is zero, the value of the second position is one, and so on. We can see this by listing each enumeration of the Day type and displaying the value of the corresponding position:

Listing 28: days.ads

```
1 package Days is
2
3 type Day is (Mon, Tue, Wed,
4 Thu, Fri,
5 Sat, Sun);
6
7 end Days;
```

Listing 29: show_days.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Days;
                      use Days;
2
3
   procedure Show_Days is
4
   begin
5
      for D in Day loop
6
         Put_Line (Day'Image (D)
7
                                        = "
                    & " position
8
                    & Integer'Image (Day'Pos (D)));
9
          Put_Line (Day'Image (D)
10
                    & " internal code = "
11
                    & Integer'Image
12
                         (Day'Enum Rep (D)));
13
       end loop;
14
   end Show_Days;
15
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Enumerations.Enumeration_Values
MD5: d6c5cb99b9770893b7277c470f40e805

Runtime output

```
MON position
                     0
                  =
MON internal code =
                     0
TUE position
                  =
                     1
TUE internal code =
                     1
WED position =
                     2
WED internal code =
                     2
THU position
                     3
                  =
THU internal code =
                     3
                     4
FRI position
                  =
FRI internal code =
                     4
SAT position
                     5
SAT internal code =
                     5
SUN position
                     6
                  =
SUN internal code =
                     6
```

Note that this application also displays the internal code, which, in this case, is equivalent to the position value for all enumerations.

We may, however, change the internal code of an enumeration using a representation clause. We discuss this topic *in another section* (page 79).

1.5 Universal and Root Types

Previously, in the section about *scalar types* (page 14), we said that scalar types are the most basic types that we can get. However, Ada has the concept of universal and root types, which could be considered *more basic* than scalar types. In fact, universal and root types are underlying scalar types used by the language designers to define the language semantics. In this section, we briefly introduce this topic.

1.5.1 Universal Types

The Ada standard defines four universal types:

- 1. universal integer types
- 2. universal real types
- 3. universal fixed types
- 4. universal access types

The first three are numeric types, and we discuss them in detail later on *in another chapter* (page 362). The last one is used for *anonymous access types* (page 711).

Universal types aren't types we can use directly, but rather via specific languages constructs. In this sense, we cannot derive from universal types, but only make use of them indirectly.

For instance, if we declare *named numbers* (page 13) using a real value, we're indirectly using a universal real type. If we declare another named number using an expression, the computation is performed based on the universal types of the elements of that expression:

```
package Show Universal Real Integer is
1
2
      Pi
3
              : constant := 3.1415926535;
                              ^^^^^
4
       - -
                         universal real type
5
       - -
6
       Two Pi : constant := Pi * 2.0:
7
                               ~~~~~
       - -
8
                            operation on
9
       - -
                        universal real type
       - -
10
11
              : constant := 10;
      Ν
12
                              ~~
       - -
13
                     universal integer type
       - -
14
15
      N 10
              : constant := N * 10;
16
                              ~~~~~
       - -
17
       - -
                           operation on
18
                     universal integer type
19
20
   end Show Universal Real Integer;
21
```

Listing 30: show_universal_real_integer.ads

Code block metadata

In this example, the expression Pi * 2.0 is computed using universal real types, while the expression N * 10 is computed using universal integer types.

Similarly, for anonymous access types, the equality operator uses universal access types for the comparison:

Listing 31: show_universal_access.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Universal Access is
3
      I : aliased Integer;
4
      A : access Integer := I'Access;
5
      B : access Integer := I'Access;
6
   begin
7
      if A = B then
8
          Put Line ("A = B");
9
      else
10
          Put Line ("A /= B");
11
      end if;
12
   end Show Universal Access;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Universal_And_Root_Types.Universal_

⇔Access

MD5: e6a37de980cc3b2c19e36baa3a51c329
```

Runtime output

A = B

In this example, both A and B are variables of anonymous access types. Because the type isn't a known named type, the equality operation = uses the universal access type for the comparison.

1 In the Ada Reference Manual

- 3.3.2 Number Declarations³¹
- 4.5.2 Relational Operators and Membership Tests³²

1.5.2 Root Types

The root types can be found on a level above the universal types. In this category, we can find the same numeric types that we have for universal types, namely the root real, root integer and root fixed types.

The term *root* is used in the context of type derivation. In fact, the root type is the first type that we derive all other types from. In other words, if we declare an integer range as a new type, that type is derived from the root integer type. Similarly, if we declare a new floating-point type, that type is derived from the root real type. For example:

Listing 32: show root integer real.ads

```
1 package Show_Root_Integer_Real is
2
3 type Score is range 0 .. 10;
4 -- Type Score is derived from
5 -- the root integer type.
```

(continues on next page)

³¹ http://www.ada-auth.org/standards/22rm/html/RM-3-3-2.html

³² http://www.ada-auth.org/standards/22rm/html/RM-4-5-2.html

```
6
7 type Real_Score is
8 digits 10 range 0.0 .. 10.0;
9 -- Type Real_Score is derived from
10 -- the root real type.
11
12 end Show_Root_Integer_Real;
```

Code block metadata

Here, Score and Real_Score are derived from the root integer and real types, respectively. Note that the derivation is always implicit, as we cannot write something like **type Score** is new Root_Integer range 0 .. 10 or **type Real_Score** is new Root_Real digits 10 range 0.0 .. 10.0.

In contrast, if we derive from an existing floating-point or integer type defined by the Ada standard, we're not deriving directly from the root types:

Listing 33: show_standard_derivation.ads

```
package Show Standard Derivation is
1
2
      type Score is new Integer
3
        range 0 .. 10;
4
      -- Type Score is derived from
5
      - -
         the Integer type.
6
7
      type Real Score is new Float
8
        range 0.0 .. 10.0;
9
       -- Type Real_Score is derived from
10
      -- the Float type.
11
12
   end Show Standard Derivation;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Universal_And_Root_Types.Standard_

⊲Integer_Float_Derivation

MD5: d32261966e1f1ae9626336f57ab16d89
```

In this case, we're explicitly deriving from the standard Ada types **Integer** and **Float**, which, on their turn, are derived from the root integer and root real types, respectively.

For further reading...

Ada also has the concept of *base types* (page 376), which *sounds* similar to the concept of the root type. However, the focus of each one is different: while the root type refers to the derivation tree of a type, the base type refers to the constraints of a type.

We discuss base types and the Base attribute (page 377) in another chapter.

1.6 Definite and Indefinite Subtypes

Indefinite types were mentioned back in the Introduction to Ada course³³. In this section, we'll recapitulate and extend on both definite and indefinite types.

Definite types are the basic kind of types we commonly use when programming applications. For example, we can only declare variables of definite types; otherwise, we get a compilation error. Interestingly, however, to be able to explain what definite types are, we need to first discuss indefinite types.

Indefinite types include:

- · unconstrained arrays;
- record types with unconstrained discriminants without defaults.

Let's see some examples of indefinite types:

```
Listing 34: unconstrained_types.ads
```

```
package Unconstrained_Types is
1
2
3
      type Integer_Array is
         array (Positive range <>) of Integer;
4
5
      type Simple Record (Extended : Boolean) is
6
      record
7
          V : Integer;
8
          case Extended is
9
             when False =>
10
                null;
11
             when True =>
12
                V_Float : Float;
13
          end case;
14
      end record;
15
16
   end Unconstrained_Types;
17
```

Code block metadata

In this example, both Integer_Array and Simple_Record are indefinite types.

As we've just mentioned, we cannot declare variable of indefinite types:

Listing 35: using_unconstrained_type.adb

```
with Unconstrained_Types; use Unconstrained_Types;
1
2
   procedure Using_Unconstrained_Type is
3
4
      A : Integer_Array;
5
6
      R : Simple Record;
7
8
   begin
9
      null:
10
   end Using_Unconstrained_Type;
11
```

Code block metadata

³³ https://learn.adacore.com/courses/intro-to-ada/chapters/arrays.html#intro-ada-indefinite-subtype

Build output

```
using_unconstrained_type.adb:5:08: error: unconstrained subtype not allowed (need

_initialization)

using_unconstrained_type.adb:5:08: error: provide initial value or explicit array

_bounds

using_unconstrained_type.adb:7:08: error: unconstrained subtype not allowed (need

_initialization)

using_unconstrained_type.adb:7:08: error: provide initial value or explicit

_discriminant values

using_unconstrained_type.adb:7:08: error: or give default discriminant values for

_type "Simple_Record"

gprbuild: *** compilation phase failed
```

As we can see when we try to build this example, the compiler complains about the declaration of A and R because we're trying to use indefinite types to declare variables. The main reason we cannot use indefinite types here is that the compiler needs to know at this point how much memory it should allocate. Therefore, we need to provide the information that is missing. In other words, we need to change the declaration so the type becomes definite. We can do this by either declaring a definite type or providing constraints in the variable declaration. For example:

Listing 36: using_unconstrained_type.adb

```
with Unconstrained_Types; use Unconstrained_Types;
1
2
   procedure Using_Unconstrained_Type is
3
4
      subtype Integer_Array_5 is
5
        Integer_Array (1 .. 5);
6
7
      A1 : Integer Array 5;
8
      A2 : Integer Array (1 .. 5);
9
10
      subtype Simple_Record_Ext is
11
         Simple_Record (Extended => True);
12
13
      R1 : Simple_Record_Ext;
14
      R2 : Simple_Record (Extended => True);
15
16
   begin
17
      null;
18
   end Using_Unconstrained_Type;
19
```

Code block metadata

In this example, we declare the Integer_Array_5 subtype, which is definite because we're constraining it to a range from 1 to 5, thereby defining the information that was missing in the indefinite type Integer_Array. Because we now have a definite type, we can use it to declare the A1 variable. Similarly, we can use the indefinite type Integer_Array directly in the declaration of A2 by specifying the previously unknown range.

Similarly, in this example, we declare the Simple_Record_Ext subtype, which is definite because we're initializing the record discriminant Extended. We can therefore use it in

the declaration of the R1 variable. Alternatively, we can simply use the indefinite type Simple_Record and specify the information required for the discriminants. This is what we do in the declaration of the R2 variable.

Although we cannot use indefinite types directly in variable declarations, they're very useful to generalize algorithms. For example, we can use them as parameters of a subprogram:

Listing 37: show_integer_array.ads

```
with Unconstrained_Types; use Unconstrained_Types;
```

```
procedure Show_Integer_Array (A : Integer_Array);
```

Listing 38: show_integer_array.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Integer_Array (A : Integer_Array)
3
   is
4
5
   begin
      for I in A'Range loop
6
         Put_Line (Positive'Image (I)
7
                    & ": "
8
                    & Integer'Image (A (I)));
9
      end loop;
10
      Put_Line ("-----");
11
   end Show_Integer_Array;
12
```



```
with Unconstrained Types; use Unconstrained Types;
1
   with Show_Integer_Array;
2
3
   procedure Using Unconstrained Type is
4
      A_5 : constant Integer_Array (1 .. 5) :=
5
                (1, 2, 3, 4, 5);
6
      A_10 : constant Integer_Array (1 .. 10) :=
7
               (1, 2, 3, 4, 5, others => 99);
8
   begin
9
      Show_Integer_Array (A_5);
10
      Show_Integer_Array (A_10);
11
   end Using Unconstrained Type;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Definite_Indefinite_Subtypes.

→Indefinite_Types

MD5: 3f744fa5921a55865bc5361ec4c6eb88
```

Runtime output

1. 1 2: 2 3: 3 4: 4 5: 5 - - - - - - - -1: 1 2: 2 3: 3 4: 4 5: 5

2

6: 99 7: 99 8: 99 9: 99 10: 99

In this particular example, the compiler doesn't know a priori which range is used for the A parameter of Show_Integer_Array. It could be a range from 1 to 5 as used for variable A_5 of the Using_Unconstrained_Type procedure, or it could be a range from 1 to 10 as used for variable A_10, or it could be anything else. Although the parameter A of Show_Integer_Array is unconstrained, both calls to Show_Integer_Array — in Us-ing_Unconstrained_Type procedure — use constrained objects.

Note that we could call the Show_Integer_Array procedure above with another unconstrained parameter. For example:

Listing 40: show integer array header.ads

```
with Unconstrained_Types; use Unconstrained_Types;
procedure Show_Integer_Array_Header
(AA : Integer_Array;
HH : String);
```

Listing 41: show_integer_array_header.adb

```
use Ada.Text_I0;
   with Ada.Text IO;
1
   with Show_Integer_Array;
2
3
   procedure Show_Integer_Array_Header
4
      (AA : Integer_Array;
5
      HH : String)
6
   is
7
   begin
8
      Put Line (HH);
9
      Show_Integer_Array (AA);
10
   end Show_Integer_Array_Header;
11
```

Listing 42: using_unconstrained_type.adb

```
with Unconstrained_Types; use Unconstrained_Types;
1
2
   with Show_Integer_Array_Header;
3
4
   procedure Using_Unconstrained_Type is
5
      A_5 : constant Integer_Array (1 .. 5) :=
6
                (1, 2, 3, 4, 5);
7
      A_10 : constant Integer_Array (1 .. 10) :=
8
                (1, 2, 3, 4, 5, others => 99);
9
   begin
10
      Show_Integer_Array_Header (A_5,
11
                                   "First example");
12
      Show_Integer_Array_Header (A_10,
13
                                   "Second example");
14
   end Using Unconstrained Type;
15
```

Code block metadata

□ Indefinite_Types MD5: dd09f8c4089c6ad4c18410879f80f731

Runtime output

First example 1: 1 2: 2 3: 3 4: 4 5: 5 Second example 1: 1 2: 2 3: 3 4: 4 5: 5 6: 99 7: 99 8: 99 9: 99 10: 99 - - - - - - - -

In this case, we're calling the Show_Integer_Array procedure with another unconstrained parameter (the AA parameter). However, although we could have a long *chain* of procedure calls using indefinite types in their parameters, we still use a (definite) object at the beginning of this chain. For example, for the A 5 object, we have this chain:

A_5

Therefore, at this specific call to Show_Integer_Array, even though A is declared as a parameter of indefinite type, the actual argument is of definite type because A_5 is constrained — and, thus, of definite type.

Note that we can declare variables based on parameters of indefinite type. For example:

Listing 43: show_integer_array_plus.ads

```
with Unconstrained_Types; use Unconstrained_Types;
procedure Show_Integer_Array_Plus
(A : Integer_Array;
V : Integer);
```

Listing 44: show_integer_array_plus.adb

```
with Show_Integer_Array;
procedure Show_Integer_Array_Plus
(A : Integer_Array;
V : Integer)
is
A_Plus : Integer_Array (A'Range);
begin
```

```
9 for I in A_Plus'Range loop
10 A_Plus (I) := A (I) + V;
11 end loop;
12 Show_Integer_Array (A_Plus);
13 end Show_Integer_Array_Plus;
```

Listing 45: using_unconstrained_type.adb

```
with Unconstrained Types; use Unconstrained Types;
1
2
   with Show_Integer_Array_Plus;
3
4
   procedure Using Unconstrained Type is
5
      A_5 : constant Integer_Array (1 .. 5) :=
6
               (1, 2, 3, 4, 5);
7
  beain
8
      Show Integer Array Plus (A 5, 5);
9
   end Using Unconstrained Type;
10
```

Code block metadata

Runtime output

1: 6 2: 7 3: 8 4: 9 5: 10

In the Show_Integer_Array_Plus procedure, we're declaring A_Plus based on the range of A, which is itself of indefinite type. However, since the object passed as an argument to Show_Integer_Array_Plus must have a constraint, A_Plus will also be constrained. For example, in the call to Show_Integer_Array_Plus using A_5 as an argument, the declaration of A_Plus becomes A_Plus : Integer_Array (1 .. 5);. Therefore, it becomes clear that the compiler needs to allocate five elements for A_Plus.

We'll see later how definite and indefinite types apply to formal parameters.

In the Ada Reference Manual

3.3 Objects and Named Numbers³⁴

1.7 Incomplete types

Incomplete types — as the name suggests — are types that have missing information in their declaration. This is a simple example:

type Incomplete;

Because this type declaration is incomplete, we need to provide the missing information at some later point. Consider the incomplete type R in the following example:

³⁴ http://www.ada-auth.org/standards/22rm/html/RM-3-3.html

Listing 46: incomplete_type_example.ads

```
package Incomplete Type Example is
1
2
      type R;
3
      -- Incomplete type declaration!
4
5
      type R is record
6
         I : Integer;
7
      end record;
8
      -- type R is now complete!
9
10
   end Incomplete_Type_Example;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Incomplete_Types.Incomplete_Types
MD5: 5ca250595f2b0cc101df286ab319982f
```

The first declaration of type R is incomplete. However, in the second declaration of R, we specify that R is a record. By providing this missing information, we're completing the type declaration of R.

It's also possible to declare an incomplete type in the private part of a package specification and its complete form in the package body. Let's rewrite the example above accordingly:

Listing 47: incomplete_type_example.ads

```
1 package Incomplete_Type_Example is
2
3 private
4
5 type R;
6 -- Incomplete type declaration!
7
8 end Incomplete_Type_Example;
```

Listing 48: incomplete_type_example.adb

```
package body Incomplete_Type_Example is

type R is record
I : Integer;
end record;
end record;
end Incomplete Type_Example;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Incomplete_Types.Incomplete_Types_2
MD5: fd2f0301b4a63887add1cb2093692ddb
```

A typical application of incomplete types is to create linked lists using *access types* (page 593) based on those incomplete types. This kind of type is called a recursive type. For example:

Listing 49: linked_list_example.ads

```
package Linked_List_Example is
```

```
type Integer_List;
3
4
      type Next is access Integer_List;
5
6
      type Integer_List is record
7
          I : Integer;
8
          N : Next;
9
      end record;
10
11
  end Linked_List_Example;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Incomplete_Types.Linked_List_Example
MD5: b2d3a048473d498bbe691bc6e38ca1e9
```

Here, the N component of Integer_List is essentially giving us access to the next element of Integer_List type. Because the Next type is both referring to the Integer_List type and being used in the declaration of the Integer_List type, we need to start with an incomplete declaration of the Integer_List type and then complete it after the declaration of Next.

Incomplete types are useful to declare *mutually dependent types* (page 177), as we'll see later on. Also, we can also have formal incomplete types, as we'll discuss later.

```
1 In the Ada Reference Manual
```

• 3.10.1 Incomplete Type Declarations³⁵

1.8 Type view

Ada distinguishes between the partial and the full view of a type. The full view is a type declaration that contains all the information needed by the compiler. For example, the following declaration of type R represents the full view of this type:

Listing 50: full_view.ads

```
1 package Full_View is
2
3 -- Full view of the R type:
4 type R is record
5 I : Integer;
6 end record;
7
8 end Full View;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Full_View
MD5: d37792287d08f9aa3d32499e233516df
```

As soon as we start applying encapsulation and information hiding — via the **private** keyword — to a specific type, we are introducing a partial view and making only that view compile-time visible to clients. Doing so requires us to introduce the private part of the package (unless already present). For example:

³⁵ http://www.ada-auth.org/standards/22rm/html/RM-3-10-1.html

```
Listing 51: partial_full_views.ads
```

```
package Partial Full Views is
1
2
       -- Partial view of the R type:
3
      type R is private;
4
5
   private
6
7
       -- Full view of the R type:
8
      type R is record
9
         I : Integer;
10
      end record;
11
12
   end Partial_Full_Views;
13
```

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Partial_Full_View
MD5: b0cf748e43b23ea6c845e283c4266ff3
```

As indicated in the example, the **type R is private** declaration is the partial view of the R type, while the **type R is record** [...] declaration in the private part of the package is the full view.

Although the partial view doesn't contain the full type declaration, it contains very important information for the users of the package where it's declared. In fact, the partial view of a private type is all that users actually need to know to effectively use this type, while the full view is only needed by the compiler.

In the previous example, the partial view indicates that R is a private type, which means that, even though users cannot directly access any information stored in this type — for example, read the value of the I component of R —, they can use the R type to declare objects. For example:

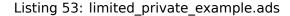
Listing 52: main.adb

```
with Partial Full Views; use Partial Full Views;
1
2
   procedure Main is
3
      -- Partial view of R indicates that
4
      -- R exists as a private type, so we
5
      -- can declare objects of this type:
6
      C : R;
7
   begin
8
      -- But we cannot directly access any
9
      -- information declared in the full
10
      - -
          view of R:
11
      - -
12
      -- C.I := 42;
13
      - -
14
      null:
15
  end Main;
16
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Partial_Full_View
MD5: 05bc9a75406d0a46f6d009d97885d010

In many cases, the restrictions applied to the partial and full views must match. For example, if we declare a limited type in the full view of a private type, its partial view must also be limited:



```
package Limited_Private_Example is
1
2
       -- Partial view must be limited,
3
       -- since the full view is limited.
4
      type R is limited private;
5
6
   private
7
8
      type R is limited record
9
         I : Integer;
10
      end record:
11
12
   end Limited Private Example;
13
```

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Limited_Private
MD5: 23d01b93fe052a500c8ca6ff76a2fd51
```

There are, however, situations where the full view may contain additional requirements that aren't mentioned in the partial view. For example, a type may be declared as non-tagged in the partial view, but, at the same time, be tagged in the full view:

Listing 54: tagged_full_view_example.ads

```
package Tagged Full View Example is
1
2
      -- Partial view using non-tagged type:
3
      type R is private;
4
5
   private
6
7
      -- Full view using tagged type:
8
      type R is tagged record
9
         I : Integer;
10
      end record;
11
12
  end Tagged_Full_View_Example;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Tagged_Full_View
MD5: 0ff9142b1ee086695b98b72a9d0f50ac
```

In this case, from a user's perspective, the R type is non-tagged, so that users cannot use any object-oriented programming features for this type. In the package body of Tagged_Full_View_Example, however, this type is tagged, so that all object-oriented programming features are available for subprograms of the package body that make use of this type. Again, the partial view of the private type contains the most important information for users that want to declare objects of this type.

In the Ada Reference Manual

• 7.3 Private Types and Private Extensions³⁶

³⁶ http://www.ada-auth.org/standards/22rm/html/RM-7-3.html

1.8.1 Non-Record Private Types

Although it's very common to declare private types as record types, this is not the only option. In fact, we could declare any type in the full view — scalars, for example —, so we could declare a "private integer" type:

Listing 55: private_integers.ads

```
package Private_Integers is
1
2
      -- Partial view of private Integer type:
3
      type Private_Integer is private;
4
5
   private
6
7
      -- Full view of private Integer type:
8
      type Private_Integer is new Integer;
9
10
   end Private Integers;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Private_Integer
MD5: flfcbed95e0f66a6f67dlbfd9ba9dflc
```

This code compiles as expected, but isn't very useful. We can improve it by adding operators to it, for example:

Listing 56: private_integers.ads

```
package Private Integers is
1
2
       -- Partial view of private Integer type:
3
      type Private Integer is private;
4
5
      function "+" (Left, Right : Private Integer)
6
                     return Private_Integer;
7
8
   private
9
10
       -- Full view of private Integer type:
11
      type Private Integer is new Integer;
12
13
   end Private Integers;
14
```

```
Listing 57: private_integers.adb
```

```
package body Private Integers is
1
2
       function "+" (Left, Right : Private_Integer)
3
                       return Private Integer
4
       is
5
          Res : constant Integer :=
6
                  Integer (Left) + Integer (Right);
7
          -- Note that we're converting Left
8
          -- and Right to Integer, which calls
9
          -- the "+" operator of the Integer
10
             type. Writing "Left + Right" would
          - -
11
          -- have called the "+" operator of
-- Private_Integer, which leads to
12
13
              recursive calls, as this is the
          - -
14
          -- operator we're currently in.
15
```

```
16 begin
17 return Private_Integer (Res);
18 end "+";
19
20 end Private_Integers;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Private_Integer
MD5: ac161cb5debfde16465c45949cf682d7
```

Now, let's use the new operator in a test application:

Listing 58: show private integers.adb

```
with Private_Integers; use Private_Integers;
procedure Show_Private_Integers is
A, B : Private_Integer;
begin
A := A + B;
end Show_Private_Integers;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Private_Integer
MD5: 5933779ce5f0802b448df96c42e65a8d
```

Build output

In this example, we use the + operator as if we were adding two common integer variables of **Integer** type.

Unconstrained Types

There are, however, some limitations: we cannot use unconstrained types such as arrays or even discriminants for arrays in the same way as we did for scalars. For example, the following declarations won't work:

Listing 59: private_arrays.ads

```
package Private Arrays is
1
2
      type Private Unconstrained Array is private;
3
4
      type Private Constrained Array
5
         (L : Positive) is private;
6
7
   private
8
9
      type Integer Array is
10
         array (Positive range <>) of Integer;
11
12
      type Private Unconstrained Array is
13
        array (Positive range <>) of Integer;
14
```

```
15
       type Private_Constrained_Array
16
         (L : Positive) is
17
           array (1 .. 2) of Integer;
18
19
          NOTE: using an array type fails as well:
20
       - -
21
       - -
           type Private Constrained Array
22
       - -
             (L : Positive) is
23
               Integer_Array (1 .. L);
24
25
   end Private_Arrays;
26
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Private_Array
MD5: b873c2d381c159532b429101e4533c05

Build output

```
private_arrays.ads:13:09: error: full view of "Private_Unconstrained_Array" not

□ compatible with declaration at line 3

private_arrays.ads:13:09: error: one is constrained, the other unconstrained

private_arrays.ads:17:07: error: elementary or array type cannot have discriminants

gprbuild: *** compilation phase failed
```

Completing the private type with an unconstrained array type in the full view is not allowed because clients could expect, according to their view, to declare objects of the type. But doing so would not be allowed according to the full view. So this is another case of the partial view having to present clients with a sufficiently *true* view of the type's capabilities.

One solution is to rewrite the declaration of **Private**_Constrained_Array using a record type:

Listing 60: private_arrays.ads

```
package Private_Arrays is
1
2
      type Private_Constrained_Array
3
         (L : Positive) is private;
4
5
   private
6
7
      type Integer_Array is
8
         array (Positive range <>) of Integer;
9
10
      type Private_Constrained_Array
11
         (L : Positive) is
12
      record
13
          Arr : Integer_Array (1 .. 2);
14
      end record;
15
16
   end Private_Arrays;
17
```

Listing 61: declare_private_array.adb

```
with Private_Arrays; use Private_Arrays;
procedure Declare_Private_Array is
Arr : Private_Constrained_Array (5);
begin
```

6 null;

7 end Declare_Private_Array;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Private_Array
MD5: 3830721499a59d85efddd4989aa7c288
```

Now, the code compiles fine — but we had to use a record type in the full view to make it work.

Another solution is to make the private type indefinite. In this case, the client's partial view would be consistent with a completion as an indefinite type in the private part:

Listing 62: private_arrays.ads

```
package Private_Arrays is
1
2
      type Private_Constrained_Array (<>) is
3
         private;
4
5
      function Init
6
         (L : Positive)
7
          return Private_Constrained_Array;
8
9
   private
10
11
      type Private Constrained Array is
12
         array (Positive range <>) of Integer;
13
14
```

15 end Private_Arrays;

Listing 63: private_arrays.adb

```
package body Private Arrays is
1
2
3
      function Init
         (L : Positive)
4
          return Private_Constrained_Array
5
6
      is
          PCA : Private_Constrained_Array (1 .. L);
7
      begin
8
          return PCA;
9
      end Init;
10
11
   end Private Arrays;
12
```

Listing 64: declare_private_array.adb

```
with Private_Arrays; use Private_Arrays;
procedure Declare_Private_Array is
Arr : Private_Constrained_Array := Init (5);
begin
null;
end Declare_Private_Array;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Type_View.Private_Array
MD5: cd170a1e44fffb93314776a68f1cb413

Build output

```
private_arrays.adb:7:07: warning: variable "PCA" is read but never assigned [-

□gnatwv]
```

The bounds for the object's declaration come from the required initial value when an object is declared. In this case, we initialize the object with a call to the Init function.

1.9 Type conversion

An important operation when dealing with objects of different types is type conversion, which we already discussed in the Introduction to Ada course³⁷. In fact, we can convert an object $0bj_X$ of an *operand* type X to a similar, closely related *target* type Y by simply indicating the target type: Y ($0bj_X$). In this section, we discuss type conversions for different kinds of types.

Ada distinguishes between two kinds of conversion: value conversion and view conversion. The main difference is the way how the operand (argument) of the conversion is evaluated:

- in a value conversion, the operand is evaluated as an *expression* (page 427);
- in a view conversion, the operand is evaluated as a name.

In other words, we cannot use expressions such as 2 * A in a view conversion, but only A. In a value conversion, we could use both forms.

In the Ada Reference Manual

```
    4.6 Type Conversions<sup>38</sup>
```

1.9.1 Value conversion

Value conversions are possible for various types. In this section, we see some examples, starting with types derived from scalar types up to array conversions.

Root and derived types

Let's start with the conversion between a scalar type and its derived types. For example, we can convert back-and-forth between the **Integer** type and the derived Int type:

Listing 65: custom_integers.ads

```
package Custom_Integers is
1
2
      type Int is new Integer
3
        with Dynamic_Predicate => Int /= 0;
4
5
      function Double (I : Integer)
6
                         return Integer is
7
         (I * 2);
8
   end Custom Integers;
10
```

³⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/strongly_typed_language.html# intro-ada-type-conversion

³⁸ http://www.ada-auth.org/standards/22rm/html/RM-4-6.html

```
Listing 66: show_conversion.adb
```

```
with Ada.Text IO;
                          use Ada.Text I0;
1
   with Custom_Integers; use Custom_Integers;
2
3
   procedure Show_Conversion is
4
      Int Var : Int
                            := 1;
5
      Integer_Var : Integer := 2;
6
   begin
7
          Int to Integer conversion
8
      Integer_Var := Integer (Int_Var);
9
10
      Put_Line ("Integer_Var : "
11
                 & Integer_Var'Image);
12
13
       -- Int to Integer conversion
14
       -- as an actual parameter
15
      Integer Var := Double (Integer (Int Var));
16
17
      Put Line ("Integer Var : "
18
                 & Integer_Var'Image);
19
20
      -- Integer to Int conversion
21
      -- using an expression
22
      Int Var
                 := Int (Integer_Var * 2);
23
24
                                ....
      Put_Line ("Int_Var :
25
                 & Int Var'Image);
26
   end Show Conversion;
27
```

Runtime output

```
Integer_Var : 1
Integer_Var : 2
Int_Var : 4
```

In the Show_Conversion procedure from this example, we first convert from Int to **Integer**. Then, we do the same conversion while providing the resulting value as an actual parameter for the Double function. Finally, we convert the Integer_Var * 2 expression from **Integer** to Int.

Note that the converted value must conform to any constraints that the target type might have. In the example above, Int has a predicate that dictates that its value cannot be zero. This (dynamic) predicate is checked at runtime, so an exception is raised if it fails:

Listing 67: show_conversion.adb

```
with Ada.Text_I0;
                         use Ada.Text_I0;
1
  with Custom_Integers; use Custom_Integers;
2
3
  procedure Show_Conversion is
4
      Int Var : Int;
5
      Integer_Var : Integer;
6
  begin
7
      Integer_Var := 0;
8
```

```
9 Int_Var := Int (Integer_Var);

10

11 Put_Line ("Int_Var : "

12 & Int_Var'Image);

13 end Show_Conversion;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Root_Derived_Type_

⊶Conversion

MD5: 4150cdffd4c1fed39fa1728a77fa599f
```

Runtime output

In this case, the conversion from **Integer** to Int fails because, while zero is a valid integer value, it doesn't obey Int's predicate.

Numeric type conversion

A typical conversion is the one between integer and floating-point values. For example:

```
Listing 68: show_conversion.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Conversion is
3
      F : Float := 1.0;
4
      I : Integer := 2;
5
   begin
6
      I := Integer (F);
7
8
      Put Line ("I : "
9
10
                  & I'Image);
11
      I := 4;
12
      F := Float (I);
13
14
      Put_Line ("F :
                         - 11
15
                 & F'Image);
16
   end Show Conversion;
17
```

Code block metadata

Runtime output

```
I : 1
F : 4.00000E+00
```

Also, we can convert between fixed-point types and floating-point or integer types:

Listing 69: fixed_point_defs.ads

```
package Fixed Point Defs is
1
      S : constant := 32;
2
      Exp : constant := 15;
3
           : constant := 2.0 ** (-S + Exp + 1);
      D
4
5
      type TQ15_31 is delta D
6
        range -1.0 * 2.0 ** Exp ..
7
               1.0 * 2.0 ** Exp - D;
8
9
      pragma Assert (TQ15_31'Size = S);
10
   end Fixed_Point_Defs;
11
```

Listing 70: show conversion.adb

```
with Fixed_Point_Defs; use Fixed_Point_Defs;
1
   with Ada.Text_I0;
                            use Ada.Text_I0;
2
3
   procedure Show_Conversion is
4
      F : Float;
5
      FP : TQ15_31;
6
      I : Integer;
7
   begin
8
9
      FP := TQ15 31 (10.25);
      I := Integer (FP);
10
11
      Put_Line ("FP : "
12
                 & FP'Image);
13
      Put_Line ("I : "
14
                 & I'Image);
15
16
      I := 128;
17
      FP := TQ15_31 (I);
18
      F := Float (FP);
19
20
      Put_Line ("FP : "
21
                 & FP'Image);
22
      Put_Line ("F :
23
                 & F'Image);
24
   end Show_Conversion;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Numeric_Type_

GConversion

MD5: 70714ba396b03469397b982e00299561
```

Runtime output

FP : 10.25000 I : 10 FP : 128.00000 F : 1.28000E+02

As we can see in the examples above, converting between different numeric types works in all directions. (Of course, rounding is applied when converting from floating-point to integer types, but this is expected.)

Enumeration conversion

We can also convert between an enumeration type and a type derived from it:

Listing 71: custom_enumerations.ads

```
package Custom_Enumerations is
type Priority is (Low, Mid, High);
type Important_Priority is new
Priority range Mid .. High;
end Custom Enumerations;
```

Listing 72: show_conversion.adb

```
with Ada.Text_I0;
                               use Ada.Text I0;
1
   with Custom_Enumerations; use Custom_Enumerations;
2
3
   procedure Show Conversion is
4
      P : Priority
                                := Low;
5
      IP : Important_Priority := High;
6
7
   begin
      P := Priority (IP);
8
9
      Put Line ("P: "
10
                 & P'Image);
11
12
      P := Mid:
13
      IP := Important_Priority (P);
14
15
      Put Line ("IP: "
16
                 & IP'Image);
17
   end Show_Conversion;
18
```

Code block metadata

Runtime output

P: HIGH IP: MID

In this example, we have the Priority type and the derived type Important_Priority. As expected, the conversion works fine when the converted value is in the range of the target type. If not, an exception is raised:

Listing 73: show_conversion.adb

```
with Ada.Text_I0;
                              use Ada.Text I0;
1
  with Custom_Enumerations; use Custom_Enumerations;
2
3
  procedure Show Conversion is
4
      P : Priority;
5
      IP : Important Priority;
6
  begin
7
      P := Low;
8
      IP := Important_Priority (P);
9
```

```
10
11 Put_Line ("IP: "
12 & IP'Image);
13 end Show_Conversion;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Enumeration_Type_ GConversion MD5: 6bbc777d4b44023bf572ca5dc6c2b4f8

Build output

Runtime output

raised CONSTRAINT_ERROR : show_conversion.adb:9 range check failed

In this example, an exception is raised because Low is not in the Important_Priority type's range.

Array conversion

Similarly, we can convert between array types. For example, if we have the array type Integer_Array and its derived type Derived_Integer_Array, we can convert between those array types:

Listing 74: custom arrays.ads

```
1 package Custom_Arrays is
2
3 type Integer_Array is
4 array (Positive range <>) of Integer;
5
6 type Derived_Integer_Array is new
7 Integer_Array;
8
9 end Custom_Arrays;
```

Listing 75: show_conversion.adb

```
with Ada.Text IO;
                         use Ada.Text I0;
1
   with Custom_Arrays; use Custom_Arrays;
2
3
   procedure Show Conversion is
4
      subtype Common Range is Positive range 1 .. 3;
5
6
      AI : Integer Array (Common Range);
7
      AI D : Derived Integer Array (Common Range);
8
   begin
9
      AI_D := [1, 2, 3];
10
      AI := Integer_Array (AI_D);
11
12
      Put Line ("AI: "
13
                 & AI'Image);
14
```

```
15

16 AI := [4, 5, 6];

17 AI_D := Derived_Integer_Array (AI);

18

19 Put_Line ("AI_D: "

20 & AI_D'Image);

21 end Show_Conversion;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Array_Type_

GConversion

MD5: 72cdf4850bec78893b6985b0c7ef02b9
```

Runtime output

AI: [1, 2, 3] AI_D: [4, 5, 6]

Note that both arrays must have the same number of components in order for the conversion to be successful. (Sliding is fine, though.) In this example, both arrays have the same range: Common_Range.

We can also convert between array types that aren't derived one from the other. As long as the components and the index subtypes are of the same type, the conversion between those types is possible. To be more precise, these are the requirements for the array conversion to be accepted:

- The component types must be the same type.
- The index types (or subtypes) must be the same or, at least, convertible.
- The dimensionality of the arrays must be the same.
- The bounds must be compatible (but not necessarily equal).

Converting between different array types can be very handy, especially when we're dealing with array types that were not declared in the same package. For example:

Listing 76: custom_arrays_1.ads

```
package Custom_Arrays_1 is
type Integer_Array_1 is
array (Positive range <>) of Integer;
type Float_Array_1 is
array (Positive range <>) of Float;
end Custom_Arrays_1;
```

Listing 77: custom_arrays_2.ads

```
1 package Custom_Arrays_2 is
2
3 type Integer_Array_2 is
4 array (Positive range <>) of Integer;
5
6 type Float_Array_2 is
7 array (Positive range <>) of Float;
```

9 end Custom_Arrays_2;

Listing 78: show_conversion.adb

```
with Ada.Text IO;
                           use Ada.Text IO;
1
   with Custom Arrays 1; use Custom Arrays 1;
2
   with Custom_Arrays_2; use Custom_Arrays_2;
3
   procedure Show_Conversion is
5
      subtype Common_Range is Positive range 1 .. 3;
6
7
      AI_1 : Integer_Array_1 (Common_Range);
8
      AI_2 : Integer_Array_2 (Common_Range);
9
      AF_1 : Float_Array_1 (Common_Range);
10
      AF_2 : Float_Array_2 (Common_Range);
11
12
   begin
      AI 2 := [1, 2, 3];
13
      AI_1 := Integer_Array_1 (AI_2);
14
15
      Put_Line ("AI_1: "
16
                 & AI_1'Image);
17
18
      AI_1 := [4, 5, 6];
19
      AI_2 := Integer_Array_2 (AI_1);
20
21
      Put Line ("AI 2: "
22
                 & AI_2'Image);
23
24
          ERROR: Cannot convert arrays whose
25
       - -
                  components have different types:
26
       - -
27
       - -
       -- AF_1 := Float_Array_1 (AI_1);
28
29
       - -
       -- Instead, use array aggregate where each
30
       -- component is converted from integer to
31
       -- float:
32
33
      AF 1 := [for I in AF 1'Range =>
34
                  Float (AI_1 (I))];
35
36
      Put_Line ("AF_1: "
37
                 & AF_1'Image);
38
39
      AF_2 := Float_Array_2 (AF_1);
40
41
      Put Line ("AF 2: "
42
                 & AF_2'Image);
43
   end Show Conversion;
44
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Array_Type_

→Conversion

MD5: 42b89fa5fe1f20af26b5da4586cea8e8
```

Runtime output

AI_1: [1, 2, 3] AI_2:

```
[ 4, 5, 6]
AF_1:
[ 4.00000E+00, 5.00000E+00, 6.00000E+00]
AF_2:
[ 4.00000E+00, 5.00000E+00, 6.00000E+00]
```

As we can see in this example, the fact that Integer_Array_1 and Integer_Array_2 have the same component type (Integer) allows us to convert between them. The same applies to the Float_Array_1 and Float_Array_2 types.

A conversion is not possible when the component types don't match. Even though we can convert between integer and floating-point types, we cannot convert an array of integers to an array of floating-point directly. Therefore, we cannot write a statement such as AF_1 := Float_Array_1 (AI_1);.

However, when the components don't match, we can of course implement the array conversion by converting the individual components. For the example above, we used an iterated component association in an array aggregate: [for I in AF_1'Range => Float (AI_1 (I))];. (We discuss this topic later *in another chapter* (page 262).)

We may also encounter array types originating from the instantiation of generic packages. In this case as well, we can use array conversions. Consider the following generic package:

```
Listing 79: custom_arrays.ads
```

```
1 generic
2 type T is private;
3 package Custom_Arrays is
4 type T_Array is
5 array (Positive range <>) of T;
6 end Custom_Arrays;
```

Code block metadata

We could instantiate this generic package and reuse parts of the previous code example:

Listing 80: show_conversion.adb

```
with Ada.Text_I0;
                         use Ada.Text_I0;
1
   with Custom_Arrays;
2
3
   procedure Show Conversion is
4
      package CA_Int_1 is
5
        new Custom_Arrays (T => Integer);
6
      package CA Int 2 is
7
        new Custom_Arrays (T => Integer);
8
9
      subtype Common_Range is Positive range 1 .. 3;
10
11
      AI 1 : CA Int 1.T Array (Common Range);
12
      AI_2 : CA_Int_2.T_Array (Common_Range);
13
   begin
14
      AI 2 := [1, 2, 3];
15
      AI_1 := CA_Int_1.T_Array (AI_2);
16
17
      Put Line ("AI 1: "
18
                 & AI_1'Image);
19
```

```
20
21 AI_1 := [4, 5, 6];
22 AI_2 := CA_Int_2.T_Array (AI_1);
23
24 Put_Line ("AI_2: "
25 & & AI_2'Image);
26 end Show_Conversion;
```

Code block metadata

Runtime output

```
AI_1:
[ 1, 2, 3]
AI_2:
[ 4, 5, 6]
```

As we can see in this example, each of the instantiated CA_Int_1 and CA_Int_2 packages has a T_Array type. Even though these T_Array types have the same name, they're actually completely unrelated types. However, we can still convert between them in the same way as we did in the previous code examples.

1.9.2 View conversion

As mentioned before, view conversions just allow names to be converted. Thus, we cannot use expressions in this case.

Note that a view conversion never changes the value during the conversion. We could say that a view conversion is simply making us *view* an object from a different angle. The object itself is still the same for both the original and the target types.

For example, consider this package:

```
Listing 81: some_tagged_types.ads
```

```
package Some_Tagged_Types is
1
2
3
       type T is tagged record
         A : Integer;
4
      end record;
5
6
      type T_Derived is new T with record
7
         B : Float:
8
      end record;
9
10
      Obj : T_Derived;
11
12
   end Some_Tagged_Types;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Tagged_Types_View
MD5: 2e18ba972682f1ae1d38e38842fde48e
```

Here, Obj is an object of type T_Derived. When we *view* this object, we notice that it has two components: A and B. However, we could *view* this object as being of type T. From that perspective, this object only has one component: A. (Note that changing the perspective

doesn't change the object itself.) Therefore, a view conversion from T_Derived to T just makes us *view* the object Obj from a different angle.

In this sense, a view conversion changes the view of a given object to the target type's view, both in terms of components that exist and operations that are available. It doesn't really change anything at all in the value itself.

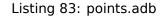
There are basically two kinds of view conversions: the ones using tagged types and the ones using untagged types. We discuss these kinds of conversion in this section.

View conversion of tagged types

A conversion between tagged types is a view conversion. Let's consider a typical code example that declares one, two and three-dimensional points:

```
Listing 82: points.ads
```

```
package Points is
1
2
3
       type Point_1D is tagged record
4
          X : Float;
      end record:
5
6
      procedure Display (P : Point_1D);
7
8
       type Point 2D is new Point 1D with record
9
         Y : Float;
10
      end record;
11
12
      procedure Display (P : Point_2D);
13
14
       type Point_3D is new Point_2D with record
15
          Z : Float;
16
      end record;
17
18
      procedure Display (P : Point 3D);
19
20
   end Points;
21
```



```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Points is
3
4
       procedure Display (P : Point 1D) is
5
       begin
6
          Put_Line ("(X => " & P.X'Image & ")");
7
       end Display;
8
9
       procedure Display (P : Point 2D) is
10
       beain
11
          Put Line ("(X => " & P.X'Image
12
                      & ", Y => " & P.Y'Image & ")");
13
       end Display;
14
15
       procedure Display (P : Point_3D) is
16
       begin
17
          Put Line ("(X => " & P.X'Image
18
                     & ", Y => " & P.Y'Image
& ", Z => " & P.Z'Image & ")");
19
20
       end Display;
21
```

22

23 end Points;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Tagged_Type_

⊶Conversion

MD5: 0acc05ae2310ab4ba038dfdb6bae0495
```

We can use the types from the Points package and convert between each other:

Listing 84: show_conversion.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Points;
                      use Points;
2
3
   procedure Show_Conversion is
4
       P_1D : Point_1D;
5
       P 3D : Point 3D;
6
   begin
7
       P 3D := (X \implies 0.1, Y \implies 0.5, Z \implies 0.3);
8
       P 1D := Point 1D (P 3D);
9
10
       Put ("P 3D : ");
11
       Display (P_3D);
12
13
       Put ("P 1D : ");
14
       Display (P_1D);
15
  end Show Conversion;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Tagged_Type_

Gonversion

MD5: fb8e07c8f2399cfae935179d8f413150
```

Runtime output

```
P_3D : (X => 1.00000E-01, Y => 5.00000E-01, Z => 3.00000E-01)
P_1D : (X => 1.00000E-01)
```

In this example, as expected, we're able to convert from the Point_3D type (which has three components) to the Point_1D type, which has only one component.

View conversion of untagged types

For untagged types, a view conversion is the one that happens when we have an object of an untagged type as an actual parameter for a formal **in out** or **out** parameter.

Let's see a code example. Consider the following simple procedure:

Listing 85: double.ads

```
procedure Double (X : in out Float);
```

Listing 86: double.adb

```
procedure Double (X : in out Float) is
begin
X := X * 2.0;
end Double;
```

The Double procedure has an **in out** parameter of **Float** type. We can call this procedure using an integer variable I as the actual parameter. For example:

Listing 87: show_conversion.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Double;
2
3
   procedure Show_Conversion is
4
       I : Integer;
5
   begin
6
       I := 2;
7
       Put_Line ("I : "
8
                   & I'Image);
9
10
       -- Calling Double with
-- Integer parameter:
11
12
       Double (Float (I));
13
       Put_Line ("I : "
14
                   & I'Image);
15
   end Show_Conversion;
16
```

Code block metadata

Runtime output

I:2 I:4

In this case, the **Float** (I) conversion in the call to Double creates a temporary floatingpoint variable. This is the same as if we had written the following code:

Listing 88: show_conversion.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Double;
2
3
   procedure Show_Conversion is
4
      I : Integer;
5
   begin
6
       I := 2;
7
      Put_Line ("I : "
8
                 & I'Image);
9
10
      declare
11
          F : Float := Float (I);
12
      begin
13
          Double (F);
14
          I := Integer (F);
15
      end;
16
      Put_Line ("I : "
17
                 & I'Image);
18
   end Show_Conversion;
19
```

Runtime output

I:2 I:4

In this sense, the view conversion that happens in Double (**Float** (I)) can be considered syntactic sugar, as it allows us to elegantly write two conversions in a single statement.

1.9.3 Implicit conversions

Implicit conversions are only possible when we have a type T and a subtype S related to the T type. For example:

Listing 89: custom_integers.ads

```
package Custom_Integers is
1
2
      type Int is new Integer
3
        with Dynamic Predicate => Int /= 0;
4
5
      subtype Sub Int 1 is Integer
6
        with Dynamic_Predicate => Sub_Int_1 /= 0;
7
8
      subtype Sub Int 2 is Sub Int 1
9
        with Dynamic_Predicate => Sub_Int_2 /= 1;
10
11
   end Custom_Integers;
12
```

Listing 90: show_conversion.adb

```
with Ada.Text_I0;
                          use Ada.Text_I0;
1
   with Custom_Integers; use Custom_Integers;
2
3
   procedure Show_Conversion is
4
      Int Var
                   : Int;
5
      Sub_Int_1_Var : Sub_Int_1;
6
      Sub_Int_2_Var : Sub_Int_2;
7
      Integer_Var : Integer;
8
   beain
9
      Integer_Var := 5;
10
                 := Int (Integer_Var);
      Int_Var
11
12
                                   n.
      Put_Line ("Int_Var :
13
                 & Int Var'Image);
14
15
      -- Implicit conversions:
16
       - -
          no explicit conversion required!
17
      Sub_Int_1_Var := Integer_Var;
18
      Sub Int 2 Var := Integer Var;
19
20
      Put Line ("Sub Int 1 Var : "
21
                 & Sub_Int_1_Var'Image);
22
      Put_Line ("Sub_Int_2_Var : "
23
                 & Sub Int 2 Var'Image);
24
   end Show_Conversion;
25
```

Runtime output

Int_Var : 5
Sub_Int_1_Var : 5
Sub_Int_2_Var : 5

In this example, we declare the Int type and the Sub_Int_1 and Sub_Int_2 subtypes:

- the Int type is derived from the Integer type,
- Sub_Int_1 is a subtype of the Integer type, and
- Sub_Int_2 is a subtype of the Sub_Int_1 subtype.

We need an explicit conversion when converting between the **Integer** and Int types. However, as the conversion is implicit for subtypes, we can simply write Sub_Int_1_Var := Integer_Var;. (Of course, writing the explicit conversion Sub_Int_1 (Integer_Var) in the assignment is possible as well.) Also, the same applies to the Sub_Int_2 subtype: we can write an implicit conversion in the Sub_Int_2_Var := Integer_Var; statement.

1.9.4 Conversion of other types

For other kinds of types, such as records, a direct conversion as we've seen so far isn't possible. In this case, we have to write a conversion function ourselves. A common convention in Ada is to name this function To_Typename. For example, if we want to convert from any type to **Integer** or **Float**, we implement the To_Integer and To_Float functions, respectively. (Obviously, because Ada supports subprogram overloading, we can have multiple To_Typename functions for different operand types.)

Let's see a code example:

Listing 91: custom_rec.ads

```
package Custom Rec is
1
2
       type Rec is record
3
         X : Integer;
4
      end record;
5
6
      function To_Integer (R : Rec)
7
                              return Integer is
8
         (R.X);
9
10
   end Custom Rec;
11
```

Listing	92:	show	conversion.adb

```
with Ada.Text IO; use Ada.Text IO;
1
  with Custom_Rec; use Custom_Rec;
2
3
   procedure Show Conversion is
4
5
      R : Rec;
      I : Integer;
6
   begin
7
      R := (X \Rightarrow 2);
8
      I := To_Integer (R);
9
```

```
10
11 Put_Line ("I : " & I'Image);
12 end Show_Conversion;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Type_Conversion.Other_Type_

GConversions

MD5: d52a4fde48243a7dd6942f0b2b91ce62
```

Runtime output

I: 2

In this example, we have the To_Integer function that converts from the Rec type to the **Integer** type.

```
In other languages
```

```
In C++, you can define conversion operators to cast between objects of different classes.
Also, you can overload the = operator. Consider this example:
#include <iostream>
class T1 {
public:
   T1 (float x) :
     x(x) {}
    // If class T3 is declared before class
    // T1, we can overload the "=" operator.
    11
    // void operator=(T3 v) {
   // x = static_cast<float>(v);
// }
   void display();
private:
  float x;
};
class T3 {
public:
   T3 (float x, float y, float z) :
     x(x), y(y), z(z) {}
    // implicit conversion
    operator float() const {
        return (x + y + z) / 3.0;
    }
    // implicit conversion
    11
    // operator T1() const {
          return T1((x + y + z) / 3.0);
    11
    11 }
    // explicit conversion (C++11)
    explicit operator T1() const {
        return T1(float(*this));
    }
```

```
void display();
private:
    float x, y, z;
};
void T1::display()
{
    std::cout << "(x => " << x</pre>
               << ")" << std::endl;
}
void T3::display()
{
    std::cout << "(x => " << x</pre>
               << "y => " << y
<< "z => " << z
               << ")" << std::endl;
}
int main ()
{
    const T3 t_3 (0.5, 0.4, 0.6);
    T1 t_1 (0.0);
    float f;
    // Implicit conversion
    f = t_3;
    std::cout << "f : " << f
               << std::endl;
    // Explicit conversion
    f = static_cast<float>(t_3);
    // f = (float)t_3;
    std::cout << "f : " << f
               << std::endl;
    // Explicit conversion
    t 1 = static cast<T1>(t 3);
    // t_1 = (T1)t_3;
    std::cout << "t_1 : ";</pre>
    t_1.display();
    std::cout << std::endl;</pre>
}
```

Here, we're using **operator float**() and **operator** T1() to cast from an object of class T3 to a floating-point value and an object of class T1, respectively. (If we switch the order and declare the T3 class before the T1 class, we could overload the = operator, as you can see in the commented-out lines.)

In Ada, this kind of conversions isn't available. Instead, we have to implement conversion functions such as the To_Integer function from the previous code example. This is the corresponding implementation:

```
Listing 93: custom_defs.ads
   package Custom_Defs is
1
2
       type T1 is private;
3
4
       function Init (X : Float)
5
6
                        return T1;
7
8
       procedure Display (Obj : T1);
9
       type T3 is private;
10
11
       function Init (X, Y, Z : Float)
12
                        return T3;
13
14
       function To_Float (Obj : T3)
15
                            return Float;
16
17
       function To_T1 (Obj : T3)
18
                         return T1;
19
20
       procedure Display (Obj : T3);
21
22
   private
23
       type T1 is record
24
          X : Float;
25
       end record;
26
27
       function Init (X : Float)
28
                        return T1 is
29
         (X => X);
30
31
       type T3 is record
32
          X, Y, Z : Float;
33
       end record;
34
35
       function Init (X, Y, Z : Float)
36
                        return T3 is
37
         (X \implies X, Y \implies Y, Z \implies Z);
38
39
   end Custom_Defs;
40
```

```
Listing 94: custom_defs.adb
   with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Custom_Defs is
3
4
       procedure Display (Obj : T1) is
5
       begin
6
          Put_Line ("(X => "
7
                     & Obj.X'Image & ")");
8
       end Display;
9
10
       function To_Float (Obj : T3)
11
                           return Float is
12
         ((0bj.X + 0bj.Y + 0bj.Z) / 3.0);
13
14
       function To_T1 (Obj : T3)
15
                        return T1 is
16
         (Init (To_Float (Obj)));
17
18
       procedure Display (Obj : T3) is
19
       begin
20
          Put Line ("(X => " & Obj.X'Image
21
                    & ", Y => " & Obj.Y'Image
22
                    & ", Z => " & Obj.Z'Image
23
                    & ")");
24
       end Display;
25
26
   end Custom Defs;
27
                               Listing 95: show conversion.adb
   with Ada.Text_IO; use Ada.Text_IO;
1
   with Custom Defs; use Custom Defs;
2
3
   procedure Show_Conversion is
4
      T_3 : constant T3 := Init (0.5, 0.4, 0.6);
5
       T_1 :
                      T1 := Init (0.0);
6
      F : Float;
7
   begin
8
      -- Explicit conversion from
9
       -- T3 to Float type
10
       F := To_Float (T_3);
11
12
       Put_Line ("F : " & F'Image);
13
14
       -- Explicit conversion from
15
       -- T3 to T1 type
16
      T 1 := To T1 (T 3);
17
18
       Put ("T_1 : ");
19
       Display (T_1);
20
   end Show Conversion;
21
```

Code block metadata

Runtime output

```
F : 5.00000E-01
T_1 : (X => 5.00000E-01)
```

In this example, we *translate* the casting operators from the C++ version by implementing the To_Float and To_T1 functions. (In addition to that, we replace the C++ constructors by Init functions.)

1.10 Qualified Expressions

We already saw qualified expressions in the Introduction to Ada³⁹ course. As mentioned there, a qualified expression specifies the exact type or subtype that the target expression will be resolved to, and it can be either any expression in parentheses, or an aggregate:

Listing 96: simple integers.ads

```
1 package Simple_Integers is
2
3 type Int is new Integer;
4
5 subtype Int_Not_Zero is Int
6 with Dynamic_Predicate => Int_Not_Zero /= 0;
7
8 end Simple_Integers;
```

Listing 97: show_qualified_expressions.adb

```
with Simple_Integers; use Simple_Integers;
procedure Show_Qualified_Expressions is
I : Int;
begin
-- Using qualified expression Int'(N)
I := Int'(0);
end Show_Qualified_Expressions;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Qualified_Expressions.Example
MD5: 0a83e10b51c72827e322984bd5c8009d
```

Here, the qualified expression Int'(0) indicates that the value zero is of Int type.

1 In the Ada Reference Manual

4.7 Qualified Expressions⁴⁰

1.10.1 Verifying subtypes

\rm 1 Note

This feature was introduced in Ada 2022.

We can use qualified expressions to verify a subtype's predicate:

³⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/more_about_types.html#intro-ada-qualified-expressions
⁴⁰ http://www.ada-auth.org/standards/22rm/html/RM-4-7.html

Listing 98: show_qualified_expressions.adb

```
with Simple_Integers; use Simple_Integers;
procedure Show_Qualified_Expressions is
I : Int;
begin
I := Int_Not_Zero'(0);
end Show_Qualified_Expressions;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Qualified_Expressions.Example
MD5: 3c4ab8ad7bf75ae029047f673aa15d70
```

Build output

Runtime output

Here, the qualified expression Int_Not_Zero'(0) checks the dynamic predicate of the subtype. (This predicate check fails at runtime.)

1.11 Default initial values

In the Introduction to Ada course⁴¹, we've seen that record components can have default values. For example:

Listing 99: defaults.ads

```
1 package Defaults is
2
3 type R is record
4 X : Positive := 1;
5 Y : Positive := 10;
6 end record;
7
8 end Defaults;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Default_Initial_Values.Defaults_1
MD5: e230be602cbb24a854e71c8176c7148c
```

In this section, we'll extend the concept of default values to other kinds of type declarations, such as scalar types and arrays.

To assign a default value for a scalar type declaration — such as an enumeration and a new integer —, we use the Default_Value aspect:

⁴¹ https://learn.adacore.com/courses/intro-to-ada/chapters/records.html#intro-ada-record-default-values

Listing 100: defaults.ads

```
package Defaults is
type E is (E1, E2, E3)
with Default_Value => E1;
type T is new Integer
with Default_Value => -1;
end Defaults;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Default_Initial_Values.Defaults_2
MD5: e6cd8261b099278ceeb5fda91d318f6e
```

Note that we cannot specify a default value for a subtype:

Listing 101: defaults.ads

```
package Defaults is
subtype T is Integer
with Default_Value => -1;
end Defaults;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Default_Initial_Values.Defaults_3
MD5: beef68e4a7a3714cfa3e547bdcda9a0c

Build output

```
defaults.ads:4:11: error: aspect "Default_Value" cannot apply to subtype
gprbuild: *** compilation phase failed
```

For array types, we use the Default_Component_Value aspect:

Listing 102: defaults.ads

```
1 package Defaults is
2
3 type Arr is
4 array (Positive range <>) of Integer
5 with Default_Component_Value => -1;
6
7 end Defaults;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Default_Initial_Values.Defaults_4
MD5: 2c390e3900e4af42498381025a37955e
```

This is a package containing the declarations we've just seen:

Listing 103: defaults.ads

```
package Defaults is
```

2

```
type E is (E1, E2, E3)
3
         with Default_Value => E1;
4
5
       type T is new Integer
6
         with Default_Value => -1;
7
8
       -- We cannot specify default
9
          values for subtypes:
       - -
10
       - -
11
          subtype T is Integer
       - -
12
       - -
             with Default Value => -1;
13
14
       type R is record
15
         X : Positive := 1;
16
         Y : Positive := 10;
17
       end record;
18
19
       type Arr is
20
         array (Positive range <>) of Integer
21
           with Default_Component_Value => -1;
22
23
   end Defaults;
24
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Default_Initial_Values.Defaults
MD5: e9263ff5b96523c129a3d2d9bbb5a4dd
```

In the example below, we declare variables of the types from the Defaults package:

Listing 104: use_defaults.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
   with Defaults; use Defaults;
2
3
   procedure Use Defaults is
4
      E1 : E;
5
      T1 : T;
6
      R1 : R;
7
      A1 : Arr (1 .. 5);
8
   begin
9
      Put_Line ("Enumeration:
10
                 & E'Image (E1));
11
      Put_Line ("Integer type: "
12
                 & T'Image (T1));
13
      Put Line ("Record type:
14
                 & Positive'Image (R1.X)
15
                 & ", "
16
                 & Positive'Image (R1.Y));
17
18
      Put ("Array type:
                            ");
19
       for V of A1 loop
20
          Put (Integer'Image (V) & " ");
21
      end loop;
22
      New Line;
23
   end Use Defaults;
24
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.Default_Initial_Values.Defaults
MD5: f8e55d31cbda2447fe14eb07eaad1975

Runtime output

Enumeration: E1 Integer type: -1 Record type: 1, 10 Array type: -1 -1 -1 -1 -1

As we see in the Use_Defaults procedure, all variables still have their default values, since we haven't assigned any value to them.

1 In the Ada Reference Manual

- 3.5 Scalar Types⁴²
- 3.6 Array Types⁴³

1.12 Deferred Constants

Deferred constants are declarations where the value of the constant is not specified immediately, but rather *deferred* to a later point. In that sense, if a constant declaration is deferred, it is actually declared twice:

- 1. in the deferred constant declaration, and
- 2. in the full constant declaration.

The simplest form of deferred constant is the one that has a full constant declaration in the private part of the package specification. For example:

```
Listing 105: deferred_constants.ads
```

```
package Deferred Constants is
1
2
      type Speed is new Long_Float;
3
4
      Light : constant Speed;
5
               ^ deferred constant declaration
      - -
6
7
   private
8
9
      Light : constant Speed := 299 792 458.0;
10
               ^ full constant declaration
11
12
   end Deferred Constants;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Deferred_Constants.Deferred_

GConstant_Private

MD5: f76e42326889f70fa7e1e216576f9771
```

Another form of deferred constant is the one that imports a constant from an external implementation — using the Import keyword. We can use this to import a constant declaration from an implementation in C. For example, we can declare the light constant in a C file:

⁴² http://www.ada-auth.org/standards/22rm/html/RM-3-5.html

⁴³ http://www.ada-auth.org/standards/22rm/html/RM-3-6.html

Listing 106: constants.c

1 double light = 299792458.0;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Deferred_Constants.Deferred_

→Constant_C

MD5: 71194a329dc5adaac3e01aff143a9943
```

Then, we can import this constant in the Deferred_Constants package:

Listing 107: deferred_constants.ads

```
package Deferred_Constants is
1
2
      type Speed is new Long Float;
3
4
      Light : constant Speed with
5
        Import, Convention => C;
6
       -- ^^^^ deferred constant
7
               declaration; imported
       - -
8
                from C file
9
10
   end Deferred Constants;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Deferred_Constants.Deferred_

Gonstant_C

MD5: 9355d194e973c6c6540485178b2259c9
```

In this case, we don't have a full declaration in the Deferred_Constants package, as the Light constant is imported from the constants.c file.

As a rule, the deferred and the full declarations should match — except, of course, for the actual value that is missing in the deferred declaration. For instance, we're not allowed to use different types in both declarations. However, we may use a subtype in the full declaration — as long as it's compatible with the type that was used in the deferred declaration. For example:

Listing 108: deferred_constants.ads

```
package Deferred_Constants is
1
2
      type Speed is new Long_Float;
3
4
      subtype Positive_Speed is
5
         Speed range 0.0 .. Speed'Last;
6
7
      Light : constant Speed;
8
               ^ deferred constant declaration
9
10
   private
11
12
      Light : constant Positive_Speed :=
13
                 299_792_458.0;
14
               ^ full constant declaration
15
       - -
                 using a subtype
       - -
16
17
   end Deferred_Constants;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Deferred_Constants.Deferred_

Gonstant_Subtype

MD5: ad6e13e30bacb6d97ccfa6c7345ffb67
```

Here, we're using the Speed type in the deferred declaration of the Light constant, but we're using the Positive_Speed subtype in the full declaration.

A useful application of deferred constants is when the value of the constant is calculated using entities not meant to be compile-time visible to clients. As such, these other entities are only visible in the private part of the package, so that's where the value of the deferred constant must be computed. For example, the full constant declaration may be computed by a call to an expression function:

Listing 109: deferred_constants.ads

```
package Deferred_Constants is
1
2
      type Speed is new Long_Float;
3
4
      Light : constant Speed;
5
               ^ deferred constant declaration
6
      - -
7
   private
8
9
      function Calculate_Light return Speed is
10
         (299_792_458.0);
11
12
      Light : constant Speed := Calculate Light;
13
               ^ full constant declaration
       - -
14
                calling a private function
       - -
15
16
   end Deferred_Constants;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.Deferred_Constants.Deferred_

GConstant_Function

MD5: f0b1a9521af31a4b48bbd54891f1c32b
```

Here, we call the Calculate_Light function — declared in the private part of the Deferred_Constants package — for the full declaration of the Light constant.

In the Ada Reference Manual

• 7.4 Deferred Constants⁴⁴

1.13 User-defined literals

🚯 Note

This feature was introduced in Ada 2022.

Any type definition has a kind of literal associated with it. For example, integer types are associated with integer literals. Therefore, we can initialize an object of integer type with

⁴⁴ http://www.ada-auth.org/standards/22rm/html/RM-7-4.html

an integer literal:

Listing 110: simple integer literal.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
  procedure Simple Integer Literal is
3
      V : Integer;
4
  begin
5
      V := 10;
6
7
      Put Line (Integer'Image (V));
8
  end Simple Integer Literal;
a
```

Code block metadata

Runtime output

10

Here, 10 is the integer literal that we use to initialize the integer variable V. Other examples of literals are real literals and string literals, as we'll see later.

When we declare an enumeration type, we limit the set of literals that we can use to initialize objects of that type:

Listing 111: simple_enumeration.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Simple Enumeration is
3
      type Activation State is (Unknown, Off, On);
4
5
      S : Activation_State;
6
   begin
7
      S := 0n;
8
      Put_Line (Activation_State'Image (S));
9
  end Simple_Enumeration;
10
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.User_Defined_Literals.Simple_ →Enumeration MD5: 075df146fcb567817dadfdb245659773

Runtime output

ON

For objects of Activation_State type, such as S, the only possible literals that we can use are Unknown, Off and On. In this sense, types have a constrained set of literals that can be used for objects of that type.

User-defined literals allow us to extend this set of literals. We could, for example, extend the type declaration of Activation_State and allow the use of integer literals for objects of that type. In this case, we need to use the Integer_Literal aspect and specify a function that implements the conversion from literals to the type we're declaring. For this conversion from integer literals to the Activation_State type, we could specify that 0 corresponds

to Off, 1 corresponds to On and other values correspond to Unknown. We'll see the corresponding implementation later.

These are the three kinds of literals and their corresponding aspect:

Literal	Example	Aspect
Integer	1	Integer_Literal
Real	1.0	Real_Literal
String	"On"	String_Literal

For our previous Activation_States type, we could declare a function Integer_To_Activation_State that converts integer literals to one of the enumeration literals that we've specified for the Activation_States type:

Listing 112: activation states.ads

```
package Activation_States is
1
2
      type Activation_State is (Unknown, Off, On)
3
        with Integer_Literal =>
4
                Integer_To_Activation_State;
5
6
      function Integer_To_Activation_State
7
         (S : String)
8
          return Activation_State;
9
10
   end Activation_States;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.User_Defined_Literals.User_Defined_ →Literals MD5: 37c497105ea3a5ad67f72955911eb31a

Based on this specification, we can now use an integer literal to initialize an object S of Activation_State type:

S : Activation_State := 1;

Note that we have a string parameter in the declaration of the Integer_To_Activation_State function, even though the function itself is only used to convert integer literals (but not string literals) to the Activation_State type. It's our job to process that string parameter in the implementation of the Integer_To_Activation_State function and convert it to an integer value — using **Integer** 'Value, for example:

Listing 113: activation_states.adb

```
package body Activation States is
1
2
      function Integer_To_Activation_State
3
        (S : String)
4
         return Activation State is
5
      begin
6
         case Integer'Value (S) is
7
                       => return Off;
            when 0
8
                        => return On;
            when 1
9
            when others => return Unknown;
10
         end case;
11
      end Integer_To_Activation_State;
12
```

```
14 end Activation_States;
```

13

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.User_Defined_Literals.User_Defined_

→Literals

MD5: c130c42ee2b91e4306c0b49bd6d5d322
```

Let's look at a complete example that makes use of all three kinds of literals:

Listing 114: activation_states.ads

```
package Activation_States is
1
2
       type Activation_State is (Unknown, Off, On)
3
         with String_Literal =>
4
                To_Activation_State,
5
              Integer_Literal =>
6
                Integer_To_Activation_State,
7
              Real_Literal
                               =>
8
                Real_To_Activation_State;
9
10
       function To_Activation_State
11
         (S : Wide_Wide_String)
12
          return Activation_State;
13
14
      function Integer_To_Activation_State
15
         (S : String)
16
          return Activation_State;
17
18
      function Real_To_Activation_State
19
         (S : String)
20
          return Activation_State;
21
22
   end Activation_States;
23
```

Listing	115:	activation	states.adb

```
package body Activation_States is
1
2
       function To_Activation_State
3
         (S : Wide_Wide_String)
4
          return Activation_State
5
       is
6
       begin
7
          if S = "Off" then
8
              return Off;
9
          elsif S = "On" then
10
              return On;
11
          else
12
              return Unknown;
13
          end if;
14
       end To_Activation_State;
15
16
       function Integer_To_Activation_State
17
         (S : String)
18
          return Activation_State
19
       is
20
       begin
21
          case Integer'Value (S) is
22
```

```
when 0
                          => return Off;
23
             when 1
                          => return On;
24
             when others => return Unknown;
25
          end case;
26
       end Integer_To_Activation_State;
27
28
       function Real_To_Activation_State
29
         (S : String)
30
          return Activation_State
31
       is
32
          V : constant Float := Float'Value (S);
33
       begin
34
          if V < 0.0 then
35
             return Unknown;
36
          elsif V < 1.0 then
37
             return Off;
38
          else
39
              return On;
40
          end if;
41
       end Real_To_Activation_State;
42
43
   end Activation_States;
44
```

Listing 116: activation_examples.adb

```
with Ada.Text I0;
                             use Ada.Text I0;
1
   with Activation_States; use Activation_States;
2
3
4
   procedure Activation_Examples is
      S : Activation_State;
5
6
   begin
      S := "Off":
7
      Put_Line ("String: Off => "
8
                 & Activation_State'Image (S));
9
10
      S := 1;
11
      Put Line ("Integer: 1
                              => "
12
                 & Activation State'Image (S));
13
14
      S := 1.5;
15
      Put_Line ("Real:
                           1.5 => "
16
                 & Activation_State'Image (S));
17
   end Activation_Examples;
18
```

Code block metadata

Runtime output

String: Off => OFF
Integer: 1 => ON
Real: 1.5 => ON

In this example, we're extending the declaration of the Activation_State type to include string and real literals. For string literals, we use the To_Activation_State function, which converts:

• the "Off" string to Off,

- the "On" string to On, and
- any other string to Unknown.

For real literals, we use the Real_To_Activation_State function, which converts:

- any negative number to Unknown,
- a value in the interval [0, 1) to 0ff, and
- a value equal or above 1.0 to 0n.

Note that the string parameter of To_Activation_State function — which converts string literals — is of Wide_Wide_String type, and not of **String** type, as it's the case for the other conversion functions.

In the Activation_Examples procedure, we show how we can initialize an object of Activation_State type with all kinds of literals (string, integer and real literals).

With the definition of the Activation_State type that we've seen in the complete example, we can initialize an object of this type with an enumeration literal or a string, as both forms are defined in the type specification:

Listing 117: using string literal.ad	Listina	117:	usina	strina	literal	.adb
--------------------------------------	---------	------	-------	--------	---------	------

```
with Ada.Text IO;
                            use Ada.Text I0;
1
   with Activation_States; use Activation_States;
2
3
   procedure Using String Literal is
4
      S1 : constant Activation State := On;
5
      S2 : constant Activation_State := "On";
6
   begin
7
      Put_Line (Activation_State'Image (S1));
8
      Put Line (Activation State'Image (S2));
9
  end Using_String_Literal;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Types.User_Defined_Literals.Activation_

→States

MD5: 975a3c56e7a938a89a617dc59c5302a7
```

Runtime output

ON ON

Note we need to be very careful when designing conversion functions. For example, the use of string literals may limit the kind of checks that we can do. Consider the following misspelling of the 0ff literal:

Listing 118: misspelling_example.adb

```
with Ada.Text IO;
                            use Ada.Text I0;
1
   with Activation States; use Activation States;
2
3
   procedure Misspelling_Example is
4
      S : constant Activation State :=
5
            Offf;
6
             ^ Error: Off is misspelled.
7
   begin
8
      Put_Line (Activation_State'Image (S));
9
   end Misspelling Example;
10
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Types.User_Defined_Literals.Activation_ ⇔States MD5: 81a8ff17e0fb8c7dce18780d8c11a6ad

Build output

```
misspelling_example.adb:6:10: error: "Offf" is undefined
misspelling_example.adb:6:10: error: possible misspelling of "Off"
gprbuild: *** compilation phase failed
```

As expected, the compiler detects this error. However, this error is accepted when using the corresponding string literal:

Listing 119: misspelling example.adb

```
with Ada.Text IO:
                            use Ada.Text IO:
1
   with Activation States; use Activation States;
2
3
   procedure Misspelling Example is
4
      S : constant Activation State :=
5
            "Offf";
6
             ^ Error: Off is misspelled.
      - -
7
  beain
8
      Put_Line (Activation_State'Image (S));
9
  end Misspelling_Example;
10
```

Code block metadata

Runtime output

UNKNOWN

Here, our implementation of To_Activation_State simply returns Unknown. In some cases, this might be exactly the behavior that we want. However, let's assume that we'd prefer better error handling instead. In this case, we could change the implementation of To_Activation_State to check all literals that we want to allow, and indicate an error otherwise — by raising an exception, for example. Alternatively, we could specify this in the preconditions of the conversion function:

```
function To_Activation_State
 (S : Wide_Wide_String)
  return Activation_State
  with Pre => S = "Off" or
        S = "On" or
        S = "Unknown";
```

In this case, the precondition explicitly indicates which string literals are allowed for the To_Activation_State type.

User-defined literals can also be used for more complex types, such as records. For example:

Listing 120: silly_records.ads

```
package Silly_Records is
type Silly is record
```

```
4 X : Integer;
5 Y : Float;
6 end record
7 with String_Literal => To_Silly;
8
9 function To_Silly (S : Wide_Wide_String)
10 return Silly;
11 end Silly_Records;
```

Listing 121: silly_records.adb

```
package body Silly_Records is
1
2
       function To_Silly (S : Wide_Wide_String)
3
                           return Silly
4
      is
5
      begin
6
          if S = "Magic" then
7
             return (X => 42, Y => 42.0);
8
          else
9
             return (X => 0, Y => 0.0);
10
          end if;
11
      end To_Silly;
12
13
   end Silly Records;
14
```

Listing 122: silly magic.adb

```
with Ada.Text IO;
                       use Ada.Text I0;
1
  with Silly_Records; use Silly_Records;
2
3
  procedure Silly Magic is
4
      R1 : Silly;
5
  begin
6
      R1 := "Magic";
7
      Put_Line (R1.X'Image & ", " & R1.Y'Image);
8
  end Silly_Magic;
9
```

Code block metadata

Runtime output

42, 4.20000E+01

In this example, when we initialize an object of Silly type with a string, its components are:

- set to 42 when using the "Magic" string; or
- simply set to zero when using any other string.

Obviously, this example isn't particularly useful. However, the goal is to show that this approach is useful for more complex types where a string literal (or a numeric literal) might simplify handling those types. Used-defined literals let you design types in ways that, otherwise, would only be possible when using a preprocessor or a domain-specific language.

1 In the Ada Reference	Manual
• 4.2.1 User-Defined Lite	erals ⁴⁵

⁴⁵ http://www.ada-auth.org/standards/22rm/html/RM-4-2-1.html

TYPES AND REPRESENTATION

2.1 Enumeration Representation Clauses

We have talked about the internal code of an enumeration *in another section* (page 26). We may change this internal code by using a representation clause, which has the following format:

The value of each code in a representation clause must be distinct. However, as you can see above, we don't need to use sequential values — the values must, however, increase for each enumeration.

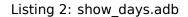
We can rewrite the previous example using a representation clause:

```
Listing 1: days.ads
```

```
2
       type Day is (Mon, Tue, Wed,
3
                     Thu, Fri,
4
                     Sat, Sun);
5
6
      for Day use (Mon => 2#00000001#,
7
                     Tue => 2#00000010#,
8
9
                     Wed => 2#00000100#,
                     Thu => 2#00001000#,
10
                     Fri => 2#00010000#,
11
                     Sat => 2#00100000#,
12
                     Sun => 2#0100000#);
13
14
   end Days;
15
```

package Days is

1



```
with Ada.Text_IO; use Ada.Text_IO;
1
   with Days;
                     use Days;
2
3
   procedure Show Days is
4
5
   begin
      for D in Day loop
6
         Put Line (Day'Image (D)
7
                    & " position
                                       = "
8
                    & Integer'Image (Day'Pos (D)));
9
         Put Line (Day'Image (D)
10
                    & " internal code = "
11
```

```
12 & Integer'Image
13 (Day'Enum_Rep (D)));
14 end loop;
15 end Show_Days;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Enumeration_

←Representation_Clauses.Enumeration_Values

MD5: a70c3f8a967c355a4bf8f2d669f9c541
```

Runtime output

MON position 0 = MON internal code = 1 TUE position = 1 TUE internal code = 2 = WED position 2 WED internal code = 4 = 3 THU position THU internal code = 8 FRI position = 4 FRI internal code = 16 SAT position = 5 SAT internal code = 32 SUN position = 6 SUN internal code = 64

Now, the value of the internal code is the one that we've specified in the representation clause instead of being equivalent to the value of the enumeration position.

In the example above, we're using binary values for each enumeration — basically viewing the integer value as a bit-field and assigning one bit for each enumeration. As long as we maintain an increasing order, we can use totally arbitrary values as well. For example:

Listing 3: days.ads

```
package Days is
1
2
       type Day is (Mon, Tue, Wed,
3
                      Thu, Fri,
4
                       Sat, Sun);
5
6
       for Day use (Mon => 5,
7
                      Tue => 9,
8
                      Wed \Rightarrow 42,
9
                      Thu => 49,
10
                      Fri => 50,
11
                      Sat => 66,
12
                      Sun => 99);
13
14
   end Days;
15
```

2.2 Data Representation

The following sections provide a glimpse on attributes and aspects used for data representation. They are usually used for embedded applications because of strict requirements that are often found there. Therefore, unless you have very specific requirements for your application, in most cases, you won't need them. However, you should at least have a rudimentary understanding of them. To read a thorough overview on this topic, please refer to the Introduction to Embedded Systems Programming⁴⁶ course.

In the Ada Reference Manual

- 13.2 Packed Types⁴⁷
- 13.3 Operational and Representation Attributes⁴⁸
- 13.5.3 Bit Ordering⁴⁹

2.3 Sizes

Ada offers multiple attributes to retrieve the size of a type or an object:

Attribute	Description
Size Object Size	Size of the representation of a subtype or an object (in bits). Size of a component or an aliased object (in bits).
Compo- nent Size	Size of a component of an array (in bits).
Storage_Size	Number of storage elements reserved for an access type or a task object.

For the first three attributes, the size is measured in bits. In the case of Storage_Size, the size is measured in storage elements. Note that the size information depends your target architecture. We'll discuss some examples to better understand the differences among those attributes.

1 Important

A storage element is the smallest element we can use to store data in memory. As we'll see soon, a storage element corresponds to a byte in many architectures.

The size of a storage element is represented by the System.Storage_Unit constant. In other words, the storage unit corresponds to the number of bits used for a single storage element.

In typical architectures, System.Storage_Unit is 8 bits. In this specific case, a storage element is equal to a byte in memory. Note, however, that System.Storage_Unit might have a value different than eight in certain architectures.

2.3.1 Size attribute and aspect

Let's start with a code example using the Size attribute:

Listing 4: custom_types.ads

package Custom_Types is

2

(continues on next page)

⁴⁶ https://learn.adacore.com/courses/intro-to-embedded-sys-prog/chapters/low_level_programming.html# intro-embedded-sys-prog-low-level-programming

⁴⁷ http://www.ada-auth.org/standards/22rm/html/RM-13-2.html

⁴⁸ http://www.ada-auth.org/standards/22rm/html/RM-13-3.html

⁴⁹ http://www.ada-auth.org/standards/22rm/html/RM-13-5-3.html

```
3 type UInt_7 is range 0 .. 127;
4 
5 type UInt_7_S32 is range 0 .. 127
6 with Size => 32;
7 
8 end Custom_Types;
```

Listing 5: show_sizes.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
    with Custom_Types; use Custom_Types;
3
4
    procedure Show_Sizes is
5
        V1 : UInt 7;
6
        V2 : UInt_7_S32;
7
    begin
8
        Put Line ("UInt 7'Size:
9
                     & UInt_7'Size'Image);
10
        Put_Line ("UInt_7'Object_Size:
                                                       n.
11
                     & UInt_7'Object_Size'Image);
12
        Put_Line ("V1'Size:
13
                     & V1'Size'Image);
14
        New Line;
15
16
        Put_Line ("UInt_7_S32'Size: "
    & UInt_7_S32'Size'Image);
Put_Line ("UInt_7_S32'Object_Size: "
    & UInt_7_S32'Object_Size'Image);
17
18
19
20
        Put_Line ("V2'Size:
21
                     & V2'Size'Image);
22
    end Show_Sizes;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Data_Representation.

⊲Sizes

MD5: e0da7cd23dc6989bea3d2902221f033e
```

Build output

```
show_sizes.adb:6:04: warning: variable "V1" is read but never assigned [-gnatwv]
show_sizes.adb:7:04: warning: variable "V2" is read but never assigned [-gnatwv]
```

Runtime output

UInt_7'Size: 7 UInt_7'Object_Size: 8 V1'Size: 8 UInt_7_S32'Size: 32 UInt_7_S32'Object_Size: 32 V2'Size: 32

Depending on your target architecture, you may see this output:

<pre>UInt_7'Size: UInt_7'Object_Size: V1'Size:</pre>	7 8 8	
<pre>UInt_7_S32'Size:</pre>	32	(continues on next page)

<pre>UInt_7_S32'Object_Size:</pre>	32
V2'Size:	32

When we use the Size attribute for a type T, we're retrieving the minimum number of bits necessary to represent objects of that type. Note that this is not the same as the actual size of an object of type T because the compiler will select an object size that is appropriate for the target architecture.

In the example above, the size of the UInt_7 is 7 bits, while the most appropriate size to store objects of this type in the memory of our target architecture is 8 bits. To be more specific, the range of UInt_7 (0 .. 127) can be perfectly represented in 7 bits. However, most target architectures don't offer 7-bit registers or 7-bit memory storage, so 8 bits is the most appropriate size in this case.

We can retrieve the size of an object of type T by using the Object_Size. Alternatively, we can use the Size attribute directly on objects of type T to retrieve their actual size — in our example, we write V1'Size to retrieve the size of V1.

In the example above, we've used both the Size attribute (for example, UInt_7'Size) and the Size aspect (with Size => 32). While the size attribute is a function that returns the size, the size aspect is a request to the compiler to verify that the expected size can be used on the target platform. You can think of this attribute as a dialog between the developer and the compiler:

(Developer) "I think that UInt_7_S32 should be stored using at least 32 bits. Do you agree?"

(Ada compiler) "For the target platform that you selected, I can confirm that this is indeed the case."

Depending on the target platform, however, the conversation might play out like this:

(Developer) "I think that UInt_7_S32 should be stored using at least 32 bits. Do you agree?"

(Ada compiler) "For the target platform that you selected, I cannot possibly do it! COMPILATION ERROR!"

2.3.2 Component size

Let's continue our discussion on sizes with an example that makes use of the Component_Size attribute:

```
Listing 6: custom_types.ads
```

```
package Custom Types is
1
2
      type UInt 7 is range 0 .. 127;
3
4
      type UInt 7 Array is
5
        array (Positive range <>) of UInt 7;
6
7
      type UInt 7 Array Comp 32 is
8
        array (Positive range <>) of UInt 7
9
           with Component Size => 32;
10
11
   end Custom_Types;
12
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Custom Types; use Custom Types;
3
4
   procedure Show Sizes is
5
       Arr_1 : UInt_7_Array (1 .. 20);
6
       Arr_2 : UInt_7_Array_Comp_32 (1 .. 20);
7
   begin
8
       Put_Line
9
                                                     н
         ("UInt_7_Array'Size:
10
          & UInt_7_Array'Size'Image);
11
       Put Line
12
         ("UInt_7_Array'Object_Size:
& UInt_7_Array'Object_Size'Image);
13
14
       Put Line
15
         ("UInt 7 Array'Component Size:
16
          & UInt_7_Array'Component_Size'Image);
17
       Put Line
18
         ("Arr_1'Component_Size:
                                                     n.
19
          & Arr_1'Component_Size'Image);
20
       Put Line
21
                                                     n
         ("Arr 1'Size:
22
          & Arr 1'Size'Image);
23
       New Line;
24
25
       Put Line
26
         ("UInt_7_Array_Comp_32'Object_Size:
27
          & UInt_7_Array_Comp_32'Size'Image);
28
       Put Line
29
         ("UInt_7_Array_Comp_32'Object_Size:
30
          & UInt_7_Array_Comp_32'Object_Size'Image);
31
       Put Line
32
         ("UInt_7_Array_Comp_32'Component_Size: "
33
          &
34
          UInt 7 Array Comp 32'Component Size'Image);
35
       Put_Line
36
         ("Arr_2'Component_Size:
                                                     n.
37
          & Arr_2'Component_Size'Image);
38
       Put Line
39
                                                     н
         ("Arr_2'Size:
40
          & Arr_2'Size'Image);
41
       New Line;
42
   end Show_Sizes;
43
```

```
Listing 7: show_sizes.adb
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Data_Representation. →Sizes MD5: e316bcb827e014075dfbf044935827ae

Build output

```
show_sizes.adb:6:04: warning: variable "Arr_1" is read but never assigned [-gnatwv]
show_sizes.adb:7:04: warning: variable "Arr_2" is read but never assigned [-gnatwv]
```

Runtime output

UInt_7_Array'Size:	17179869176	
<pre>UInt_7_Array'Object_Size:</pre>	17179869176	
<pre>UInt_7_Array'Component_Size:</pre>	8	
<pre>Arr_1'Component_Size:</pre>	8	
		(continuos on novt nago)

Arr_1'Size:	160
<pre>UInt_7_Array_Comp_32'Object_Size:</pre>	68719476704
UInt_7_Array_Comp_32'Object_Size:	68719476704
UInt_7_Array_Comp_32'Component_Size:	32
Arr_2'Component_Size:	32
Arr_2'Size:	640

Depending on your target architecture, you may see this output:

<pre>UInt_7_Array'Size:</pre>	17179869176
UInt_7_Array'Object_Size:	17179869176
UInt_7_Array'Component_Size:	8
Arr_1'Component_Size:	8
Arr_1'Size:	160
<pre>UInt_7_Array_Comp_32'Size:</pre>	68719476704
UInt_7_Array_Comp_32'Object_Size:	68719476704
UInt_7_Array_Comp_32'Component_Size:	32
Arr_2'Component_Size:	32
Arr_2'Size:	640

Here, the value we get for Component_Size of the UInt_7_Array type is 8 bits, which matches the UInt_7'Object_Size — as we've seen in the previous subsection. In general, we expect the component size to match the object size of the underlying type.

However, we might have component sizes that aren't equal to the object size of the component's type. For example, in the declaration of the UInt_7_Array_Comp_32 type, we're using the Component_Size aspect to query whether the size of each component can be 32 bits:

type UInt_7_Array_Comp_32 is array (Positive range <>) of UInt_7 with Component_Size => 32;

If the code compiles, we see this value when we use the Component_Size attribute. In this case, even though UInt_7'Object_Size is 8 bits, the component size of the array type (UInt_7_Array_Comp_32'Component_Size) is 32 bits.

Note that we can use the Component_Size attribute with data types, as well as with actual objects of that data type. Therefore, we can write UInt_7_Array'Component_Size and Arr_1'Component_Size, for example.

This big number (17179869176 bits) for UInt_7_Array'Size and UInt_7_Array'Object_Size might be surprising for you. This is due to the fact that Ada is reporting the size of the UInt_7_Array type for the case when the complete range is used. Considering that we specified a positive range in the declaration of the UInt_7_Array type, the maximum length on this machine is $2^{31} - 1$. The object size of an array type is calculated by multiplying the maximum length by the component size. Therefore, the object size of the UInt_7_Array type corresponds to the multiplication of $2^{31} - 1$ components (maximum length) by 8 bits (component size).

2.3.3 Storage size

To complete our discussion on sizes, let's look at this example of storage sizes:

Listing 8: custom_types.ads

```
package Custom_Types is
type UInt_7 is range 0 .. 127;
type UInt_7_Access is access UInt_7;
end Custom_Types;
```

Listing 9: show sizes.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
   with System;
2
3
   with Custom_Types; use Custom_Types;
4
5
   procedure Show_Sizes is
6
      AV1, AV2 : UInt_7_Access;
7
8
   begin
9
      Put_Line
         ("UInt_7_Access'Storage_Size:
10
          & UInt_7_Access'Storage_Size'Image);
11
      Put Line
12
         ("UInt_7_Access'Storage_Size (bits):
13
          & Integer'Image (UInt_7_Access'Storage_Size
14
                            * System.Storage_Unit));
15
16
      Put_Line
17
                                         n.
         ("UInt_7'Size:
18
          & UInt_7'Size'Image);
19
      Put_Line
20
         ("UInt 7 Access'Size:
                                         н
21
         & UInt_7_Access'Size'Image);
22
      Put Line
23
         ("UInt 7 Access'Object Size: "
24
          & UInt_7_Access'Object_Size'Image);
25
      Put Line
26
         ("AV1'Size:
27
          & AV1'Size'Image);
28
      New_Line;
29
30
      Put_Line ("Allocating AV1...");
31
      AV1 := new UInt_7;
32
      Put_Line ("Allocating AV2...");
33
      AV2 := new UInt_7;
34
      New_Line;
35
36
      Put Line
37
                                         н
         ("AV1.all'Size:
38
          & AV1.all'Size'Image);
39
      New_Line;
40
   end Show_Sizes;
41
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Data_Representation.

⇔Sizes

MD5: 5e652ee25b8550ac331f3ce98e24f7ba
```

Runtime output

UInt_7_Access'Storage_Size: UInt_7_Access'Storage_Size (bits): UInt_7'Size: 7 UInt_7_Access'Size: 64 UInt_7_Access'Object_Size: 64 AV1'Size: 64 Allocating AV1... Allocating AV2... AV1.all'Size: 8

Depending on your target architecture, you may see this output:

0

0

0

0

```
UInt_7_Access'Storage_Size:
UInt_7_Access'Storage_Size (bits):
UInt_7'Size: 7
UInt_7_Access'Size: 64
UInt_7_Access'Object_Size: 64
AV1'Size: 64
Allocating AV1...
Allocating AV2...
AV1.all'Size: 8
```

As we've mentioned earlier on, Storage_Size corresponds to the number of storage elements reserved for an access type or a task object. In this case, we see that the storage size of the UInt_7_Access type is zero. This is because we haven't indicated that memory should be reserved for this data type. Thus, the compiler doesn't reserve memory and simply sets the size to zero.

Because Storage_Size gives us the number of storage elements, we have to multiply this value by System.Storage_Unit to get the total storage size in bits. (In this particular example, however, the multiplication doesn't make any difference, as the number of storage elements is zero.)

Note that the size of our original data type UInt_7 is 7 bits, while the size of its corresponding access type UInt_7_Access (and the access object AV1) is 64 bits. This is due to the fact that the access type doesn't contain an object, but rather memory information about an object. You can retrieve the size of an object allocated via **new** by first dereferencing it — in our example, we do this by writing AV1.**all**'Size.

Now, let's use the Storage_Size aspect to actually reserve memory for this data type:

Listing 10: custom_types.ads

```
1 package Custom_Types is
2
3 type UInt_7 is range 0 .. 127;
4
5 type UInt_7_Reserved_Access is access UInt_7
6 with Storage_Size => 8;
7
8 end Custom_Types;
```

Listing 11: show_sizes.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with System;
```

```
3
   with Custom_Types; use Custom_Types;
4
5
   procedure Show_Sizes is
6
      RAV1, RAV2 : UInt_7_Reserved_Access;
7
   begin
8
      Put_Line
9
       ("UInt_7_Reserved_Access'Storage_Size:
10
       & UInt_7_Reserved_Access'Storage_Size'Image);
11
12
      Put Line
13
       ("UInt_7_Reserved_Access'Storage_Size (bits): "
14
15
        & Integer'Image
            (UInt_7_Reserved_Access'Storage_Size
16
             * System.Storage_Unit));
17
18
      Put Line
19
         ("UInt 7 Reserved Access'Size:
20
         & UInt_7_Reserved_Access'Size'Image);
21
      Put Line
22
         ("UInt 7 Reserved Access'Object Size: "
23
         & UInt_7_Reserved_Access'Object_Size'Image);
24
      Put Line
25
                                                  n.
         ("RAV1'Size:
26
         & RAV1'Size'Image);
27
      New_Line;
28
29
      Put_Line ("Allocating RAV1...");
30
      RAV1 := new UInt_7;
31
      Put_Line ("Allocating RAV2...");
32
      RAV2 := new UInt 7;
33
      New Line;
34
   end Show_Sizes;
35
```

Code block metadata

Runtime output

```
UInt_7_Reserved_Access'Storage_Size: 8
UInt_7_Reserved_Access'Storage_Size (bits): 64
UInt_7_Reserved_Access'Size: 64
RAV1'Size: 64
Allocating RAV1...
Allocating RAV2...
raised STORAGE_ERROR : s-poosiz.adb:108 explicit raise
```

Depending on your target architecture, you may see this output:

```
UInt_7_Reserved_Access'Storage_Size: 8
UInt_7_Reserved_Access'Storage_Size (bits): 64
UInt_7_Reserved_Access'Size: 64
UInt_7_Reserved_Access'Object_Size: 64
RAV1'Size: 64
```

```
Allocating RAV1...
Allocating RAV2...
raised STORAGE_ERROR : s-poosiz.adb:108 explicit raise
```

In this case, we're reserving 8 storage elements in the declaration of UInt 7 Reserved Access.

```
type UInt_7_Reserved_Access is access UInt_7
with Storage_Size => 8;
```

Since each storage element corresponds to one byte (8 bits) in this architecture, we're reserving a maximum of 64 bits (or 8 bytes) for the UInt_7_Reserved_Access type.

This example raises an exception at runtime — a storage error, to be more specific. This is because the maximum reserved size is 64 bits, and the size of a single access object is 64 bits as well. Therefore, after the first allocation, the reserved storage space is already consumed, so we cannot allocate a second access object.

This behavior might be quite limiting in many cases. However, for certain applications where memory is very constrained, this might be exactly what we want to see. For example, having an exception being raised when the allocated memory for this data type has reached its limit might allow the application to have enough memory to at least handle the exception gracefully.

2.4 Alignment

For many algorithms, it's important to ensure that we're using the appropriate alignment. This can be done by using the Alignment attribute and the Alignment aspect. Let's look at this example:

```
Listing 12: custom_types.ads
```

```
package Custom_Types is
type UInt_7 is range 0 .. 127;
type Aligned_UInt_7 is new UInt_7
with Alignment => 4;
end Custom_Types;
```

```
Listing 13: show_alignment.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Custom_Types; use Custom_Types;
3
4
   procedure Show Alignment is
5
                : constant UInt 7
                                             := 0;
6
      Aligned_V : constant Aligned_UInt_7 := 0;
7
   begin
8
      Put Line
9
                                         n
         ("UInt 7'Alignment:
10
         & UInt 7'Alignment'Image);
11
      Put Line
12
         ("UInt_7'Size:
                                         n
13
          & UInt_7'Size'Image);
14
```

```
Put_Line
15
         ("UInt_7'Object_Size:
16
          & UInt_7'Object_Size'Image);
17
      Put_Line
18
                                          n
         ("V'Alignment:
19
          & V'Alignment'Image);
20
      Put_Line
21
         ("V'Size:
                                          n
22
         & V'Size'Image);
23
      New_Line;
24
25
      Put_Line
26
         ("Aligned_UInt_7'Alignment:
                                          . 11
27
          & Aligned_UInt_7'Alignment'Image);
28
      Put_Line
29
         ("Aligned_UInt_7'Size:
                                          n
30
          & Aligned_UInt_7'Size'Image);
31
      Put Line
32
         ("Aligned UInt 7'Object Size: "
33
          & Aligned_UInt_7'Object_Size'Image);
34
      Put Line
35
                                          n
         ("Aligned_V'Alignment:
36
          & Aligned_V'Alignment'Image);
37
      Put_Line
38
         ("Aligned_V'Size:
39
          & Aligned_V'Size'Image);
40
      New_Line;
41
   end Show_Alignment;
42
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Data_Representation. ⇔Alignment MD5: a2fea340559193c293ccaee226de2558

Runtime output

UInt_7'Alignment:	1
UInt_7'Size:	7
<pre>UInt_7'Object_Size:</pre>	8
V'Alignment:	1
V'Size:	8
Aligned_UInt_7'Alignment: Aligned_UInt_7'Size:	4 7
Aligned_UInt_7'Object_Size:	32
Aligned_V'Alignment:	4
Aligned_V'Size:	32

Depending on your target architecture, you may see this output:

```
UInt 7'Alignment:
                              1
UInt_7'Size:
                              7
UInt_7'Object_Size:
                              8
V'Alignment:
                              1
V'Size:
                              8
Aligned_UInt_7'Alignment:
                              4
Aligned_UInt_7'Size:
                              7
Aligned_UInt_7'Object_Size:
                              32
```

Aligned_V'Alignment:	4
Aligned_V'Size:	32

In this example, we're reusing the UInt_7 type that we've already been using in previous examples. Because we haven't specified any alignment for the UInt_7 type, it has an alignment of 1 storage unit (or 8 bits). However, in the declaration of the Aligned_UInt_7 type, we're using the Alignment aspect to request an alignment of 4 storage units (or 32 bits):

```
type Aligned_UInt_7 is new UInt_7
with Alignment => 4;
```

When using the Alignment attribute for the Aligned_UInt_7 type, we can confirm that its alignment is indeed 4 storage units (bytes).

Note that we can use the Alignment attribute for both data types and objects — in the code above, we're using UInt_7'Alignment and V'Alignment, for example.

Because of the alignment we're specifying for the Aligned_UInt_7 type, its size — indicated by the Object_Size attribute — is 32 bits instead of 8 bits as for the UInt_7 type.

Note that you can also retrieve the alignment associated with a class using S'Class'Alignment. For example:

Listing 14:	show	class	alignment.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show Class Alignment is
3
4
      type Point 1D is tagged record
5
         X : Integer;
6
      end record;
7
8
      type Point_2D is new Point_1D with record
9
         Y : Integer;
10
      end record
11
        with Alignment => 16;
12
13
      type Point 3D is new Point 2D with record
14
         Z : Integer;
15
      end record;
16
17
   begin
18
      Put_Line ("1D_Point'Alignment:
19
                 & Point_1D'Alignment'Image);
20
      Put_Line ("1D_Point'Class'Alignment:
21
                 & Point_1D'Class'Alignment'Image);
22
      Put_Line ("2D_Point'Alignment:
23
                 & Point 2D'Alignment'Image);
24
      Put_Line ("2D_Point'Class'Alignment: "
25
                 & Point 2D'Class'Alignment'Image);
26
      Put_Line ("3D_Point'Alignment:
27
                 & Point_3D'Alignment'Image);
28
      Put_Line ("3D_Point'Class'Alignment: '
29
                 & Point_3D'Class'Alignment'Image);
30
31
   end Show_Class_Alignment;
```

2.5 Overlapping Storage

Algorithms can be designed to perform in-place or out-of-place processing. In other words, they can take advantage of the fact that input and output arrays share the same storage space or not.

We can use the Has_Same_Storage and the Overlaps_Storage attributes to retrieve more information about how the storage space of two objects related to each other:

- the Has_Same_Storage attribute indicates whether two objects have the exact same storage.
 - A typical example is when both objects are exactly the same, so they obviously share the same storage. For example, for array A, A'Has_Same_Storage (A) is always True.
- the Overlaps_Storage attribute indicates whether two objects have at least one bit in common.
 - Note that, if two objects have the same storage, this implies that their storage also overlaps. In other words, A'Has_Same_Storage (B) = True implies that A'Overlaps_Storage (B) = True.

Let's look at this example:

Listing 15: int array processing.ads

```
package Int Array Processing is
1
2
      type Int_Array is
3
        array (Positive range <>) of Integer;
4
5
      procedure Show_Storage (X : Int_Array;
6
                                Y : Int_Array);
7
8
      procedure Process (X :
                                 Int Array;
9
                           Y : out Int Array);
10
11
   end Int_Array_Processing;
12
```

Listing 16: int_array_processing.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Int Array Processing is
3
4
      procedure Show Storage (X : Int Array;
5
                                 Y : Int Array) is
6
      beain
7
          if X'Has Same Storage (Y) then
8
             Put Line
9
             ("Info: X and Y have the same storage.");
10
          else
11
             Put Line
12
               ("Info: X and Y don't have"
13
                & "the same storage.");
14
          end if;
15
          if X'Overlaps Storage (Y) then
16
             Put Line
17
               ("Info: X and Y overlap.");
18
          else
19
             Put Line
20
               ("Info: X and Y don't overlap.");
21
```

```
end if;
22
       end Show_Storage;
23
24
       procedure Process (X : Int_Array;
25
                            Y : out Int_Array) is
26
       begin
27
          Put_Line ("==== PROCESS =====");
28
          Show_Storage (X, Y);
29
30
          if X'Has_Same_Storage (Y) then
31
             Put_Line ("In-place processing...");
32
          else
33
              if not X'Overlaps_Storage (Y) then
34
                 Put_Line
35
                   ("Out-of-place processing...");
36
             else
37
                 Put Line
38
                   ("Cannot process "
39
                    & "overlapping arrays...");
40
             end if;
41
          end if;
42
          New_Line;
43
       end Process;
44
45
   end Int_Array_Processing;
46
```

```
Listing 17: main.adb
```

```
with Int_Array_Processing;
1
   use Int_Array_Processing;
2
3
   procedure Main is
4
      A : Int_Array (1 .. 20) := (others => 3);
5
      B : Int_Array (1 .. 20) := (others => 4);
6
   begin
7
      Process (A, A);
8
      -- In-place processing:
9
      -- sharing the exact same storage
10
11
      Process (A (1 .. 10), A (10 .. 20));
12
      -- Overlapping one component: A (10)
13
14
      Process (A (1 .. 10), A (11 .. 20));
15
      -- Out-of-place processing:
16
      -- same array, but not sharing any storage
17
18
      Process (A, B);
19
      -- Out-of-place processing:
20
      - -
          two different arrays
21
   end Main;
22
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Data_Representation.

Goverlapping_Storage

MD5: 0f599163c6f24c3ef46ec6577b501c21
```

Build output

Runtime output

```
==== PROCESS ====
Info: X and Y have the same storage.
Info: X and Y overlap.
In-place processing...
==== PROCESS ====
Info: X and Y don't have he same storage.
Info: X and Y overlap.
Cannot process overlapping arrays...
==== PROCESS ====
Info: X and Y don't have he same storage.
Info: X and Y don't overlap.
Out-of-place processing...
==== PROCESS ====
Info: X and Y don't have the same storage.
Info: X and Y don't overlap.
Out-of-place processing...
```

In this code example, we implement two procedures:

- Show_Storage, which shows storage information about two arrays by using the Has_Same_Storage and Overlaps_Storage attributes.
- Process, which are supposed to process an input array X and store the processed data in the output array Y.
 - Note that the implementation of this procedure is actually just a mock-up, so that no processing is actually taking place.

We have four different instances of how we can call the Process procedure:

- in the Process (A, A) call, we're using the same array for the input and output arrays. This is a perfect example of in-place processing. Because the input and the output arrays arguments are actually the same object, they obviously share the exact same storage.
- in the Process (A (1 .. 10), A (10 .. 20)) call, we're using two slices of the A array as input and output arguments. In this case, a single component of the A array is shared: A (10). Because the storage space is overlapping, but not exactly the same, neither in-place nor out-of-place processing can usually be used in this case.
- in the Process (A (1 .. 10), A (11 .. 20)) call, even though we're using the same array A for the input and output arguments, we're using slices that are completely independent from each other, so that the input and output arrays are not sharing any storage in this case. Therefore, we can use out-of-place processing.
- in the Process (A, B) call, we have two different arrays which obviously don't share any storage space —, so we can use out-of-place processing.

2.6 Packed Representation

As we've seen previously, the minimum number of bits required to represent a data type might be less than the actual number of bits used to store an object of that same type. We've seen an example where UInt_7'Size was 7 bits, while UInt_7'Object_Size was 8 bits. The most extreme case is the one for the Boolean type: in this case, Boolean'Size is 1 bit, while Boolean'Object_Size might be 8 bits (or even more on certain architectures). In such cases, we have 7 (or more) unused bits in memory for each object of Boolean type. In other words, we're wasting memory. On the other hand, we're gaining speed of access

because we can directly access each element without having to first change its internal representation back and forth. We'll come back to this point later.

The situation is even worse when implementing bit-fields, which can be declared as an array of **Boolean** components. For example:

Listing 18: flag_definitions.ads

```
1 package Flag_Definitions is
2
3 type Flags is
4 array (Positive range <>) of Boolean;
5
6 end Flag Definitions;
```

Listing 19: show_flags.adb

```
with Ada.Text IO;
                            use Ada.Text I0;
1
   with Flag_Definitions; use Flag_Definitions;
2
3
   procedure Show Flags is
4
      Flags_1 : Flags (1 .. 8);
5
   begin
6
      Put_Line ("Boolean'Size:
7
                 & Boolean 'Size'Image);
8
      Put_Line ("Boolean'Object_Size:
9
                 & Boolean'Object_Size'Image);
10
      Put_Line ("Flags_1'Size:
11
                 & Flags_1'Size'Image);
12
      Put_Line ("Flags_1'Component Size: "
13
                 & Flags 1'Component Size'Image);
14
   end Show_Flags;
15
```

Code block metadata

Build output

Runtime output

Boolean'Size: 1 Boolean'Object_Size: 8 Flags_1'Size: 64 Flags_1'Component_Size: 8

Depending on your target architecture, you may see this output:

Boolean'Size: 1 Boolean'Object_Size: 8 Flags_1'Size: 64 Flags_1'Component_Size: 8

In this example, we're declaring the Flags type as an array of **Boolean** components. As we can see in this case, although the size of the **Boolean** type is just 1 bit, an object of this type has a size of 8 bits. Consequently, each component of the Flags type has a size of 8 bits. Moreover, an array with 8 components of **Boolean** type — such as the Flags_1 array — has a size of 64 bits.

Therefore, having a way to compact the representation — so that we can store multiple objects without wasting storage space — may help us improving memory usage. This is actually possible by using the Pack aspect. For example, we could extend the previous example and declare a Packed_Flags type that makes use of this aspect:

Listing 20: flag_definitions.ads

```
1
   package Flag_Definitions is
2
      type Flags is
3
        array (Positive range <>) of Boolean;
4
5
      type Packed Flags is
6
        array (Positive range <>) of Boolean
7
          with Pack;
8
9
   end Flag_Definitions;
10
```

Listing 21: show packed flags.adb

```
with Ada.Text IO;
                           use Ada.Text I0;
1
   with Flag Definitions; use Flag Definitions;
2
3
   procedure Show Packed Flags is
4
      Flags 1 : Flags (1 .. 8);
5
      Flags_2 : Packed_Flags (1 .. 8);
6
   begin
7
      Put Line ("Boolean'Size:
                                            н
8
                 & Boolean 'Size'Image);
9
      Put_Line ("Boolean'Object Size:
10
                 & Boolean'Object_Size'Image);
11
      Put Line ("Flags 1'Size:
12
                 & Flags 1'Size'Image);
13
      Put_Line ("Flags_1'Component_Size: "
14
                 & Flags_1'Component_Size'Image);
15
      Put_Line ("Flags_2'Size:
16
                 & Flags_2'Size'Image);
17
      Put_Line ("Flags_2'Component_Size: "
18
                 & Flags_2'Component_Size'Image);
19
   end Show Packed Flags;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Data_Representation.

→Packed_Flags

MD5: c71cf68dc8bc41d0df2a5e3eb61b51fd
```

Build output

Runtime output

```
Boolean'Size: 1
Boolean'Object_Size: 8
Flags_1'Size: 64
Flags_1'Component_Size: 8
Flags_2'Size: 8
Flags 2'Component Size: 1
```

Depending on your target architecture, you may see this output:

```
Boolean'Size: 1
Boolean'Object_Size: 8
Flags_1'Size: 64
Flags_1'Component_Size: 8
Flags_2'Size: 8
Flags_2'Component_Size: 1
```

In this example, we're declaring the Flags_2 array of Packed_Flags type. Its size is 8 bits — instead of the 64 bits required for the Flags_1 array. Because the array type Packed_Flags is packed, we can now effectively use this type to store an object of **Boolean** type using just 1 bit of the memory, as indicated by the Flags_2'Component_Size attribute.

In many cases, we need to convert between a *normal* representation (such as the one used for the Flags_1 array above) to a packed representation (such as the one for the Flags_2 array). In many programming languages, this conversion may require writing custom code with manual bit-shifting and bit-masking to get the proper target representation. In Ada, however, we just need to indicate the actual type conversion, and the compiler takes care of generating code containing bit-shifting and bit-masking to performs the type conversion.

Let's modify the previous example and introduce this type conversion:

Listing 22: flag_definitions.ads

```
package Flag Definitions is
1
2
      type Flags is
3
         array (Positive range <>) of Boolean;
4
5
      type Packed Flags is
6
         array (Positive range <>) of Boolean
7
          with Pack;
8
9
      Default_Flags : constant Flags :=
10
         (True, True, False, True,
11
         False, False, True, True);
12
13
   end Flag Definitions;
14
```

Listing 23: show flag conversion.adb

```
with Ada.Text I0;
                            use Ada.Text I0;
1
   with Flag_Definitions; use Flag_Definitions;
2
3
   procedure Show Flag Conversion is
4
      Flags_1 : Flags (1 .. 8);
5
      Flags_2 : Packed_Flags (1 .. 8);
6
   begin
7
      Flags_1 := Default_Flags;
8
      Flags_2 := Packed_Flags (Flags_1);
9
10
      for I in Flags_2'Range loop
11
          Put Line (I'Image & ":
12
                    & Flags_1 (I)'Image & ", "
13
                    & Flags 2 (I) 'Image);
14
       end loop;
15
   end Show Flag Conversion;
16
```

Code block metadata

Generation Generation Science Science

Runtime output

1: TRUE, TRUE 2: TRUE, TRUE 3: FALSE, FALSE 4: TRUE, TRUE 5: FALSE, FALSE 6: FALSE, FALSE 7: TRUE, TRUE 8: TRUE, TRUE

In this extended example, we're now declaring Default_Flags as an array of constant flags, which we use to initialize Flags_1.

The actual conversion happens with Flags_2 := Packed_Flags (Flags_1). Here, the type conversion Packed_Flags() indicates that we're converting from the normal representation (used for the Flags type) to the packed representation (used for Packed_Flags type). We don't need to write more code than that to perform the correct type conversion.

Also, by using the same strategy, we could read information from a packed representation. For example:

```
Flags_1 := Flags (Flags_2);
```

In this case, we use Flags() to convert from a packed representation to the normal representation.

We elaborate on the topic of converting between data representations in the section on *changing data representation* (page 107).

2.6.1 Trade-offs

As indicated previously, when we're using a packed representation (vs. using a standard *unpacked* representation), we're trading off speed of access for less memory consumption. The following table summarizes this:

Representation	More speed of access	Less memory consumption
Unpacked	Х	
Packed		Х

On one hand, we have better memory usage when we apply packed representations because we may save many bits for each object. On the other hand, there's a cost associated with accessing those packed objects because they need to be unpacked before we can actually access them. In fact, the compiler generates code — using bit-shifting and bitmasking — that converts a packed representation into an unpacked representation, which we can then access. Also, when storing a packed object, the compiler generates code that converts the unpacked representation of the object into the packed representation.

This packing and unpacking mechanism has a performance cost associated with it, which results in less speed of access for packed objects. As usual in those circumstances, before using packed representation, we should assess whether memory constraints are more important than speed in our target architecture.

2.7 Record Representation and storage clauses

In this section, we discuss how to use record representation clauses to specify how a record is represented in memory. Our goal is to provide a brief introduction into the topic. If you're interested in more details, you can find a thorough discussion about record representation clauses in the Introduction to Embedded Systems Programming⁵⁰ course.

Let's start with the simple approach of declaring a record type without providing further information. In this case, we're basically asking the compiler to select a reasonable representation for that record in the memory of our target architecture.

Let's see a simple example:

Listing 24: p.ads

```
1 package P is
2
3 type R is record
4 A : Integer;
5 B : Integer;
6 end record;
7
8 end P;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_ ⇔Storage_Clauses.Rep_Clauses_1 MD5: 88171257118810bb7e02cea60ffb1ad9

Considering a typical 64-bit PC architecture with 8-bit storage units, and **Integer** defined as a 32-bit type, we get this memory representation:

position	0	1	2	3	4	5	6	7
component		A	ł			F	3	

Each storage unit is a position in memory. In the graph above, the numbers on the top (0, 1, 2, ...) represent those positions for record R.

In addition, we can show the bits that are used for components A and B:

position	0	1	2	3	4	5	6	7
bits	#07	#8#15	#16#23	#24#31	#07	#8#15	#16#23	#24#31
component	A						В	

The memory representation we see in the graph above can be described in Ada using representation clauses, as you can see in the code starting at the **for** R **use** record line in the code example below — we'll discuss the syntax and further details right after this example.

⁵⁰ https://learn.adacore.com/courses/intro-to-embedded-sys-prog/chapters/low_level_programming.html# intro-embedded-sys-prog-low-level-programming

```
Listing 25: p.ads
```

```
package P is
1
2
      type R is record
3
         A : Integer;
4
         B : Integer:
5
      end record;
6
7
       -- Representation clause for record R:
8
      for R use record
9
         A at 0 range 0 .. 31;
10
         -- ^ starting memory position
11
         B at 4 range 0 .. 31;
12
                       ^ first bit .. last bit
         - -
13
      end record;
14
15
   end P;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_

→Storage_Clauses.Rep_Clauses_2

MD5: b6be86ae7e1a5c2e7d981fe37bad49ed
```

Here, we're specifying that the A component is stored in the bits #0 up to #31 starting at position #0. Note that the position itself doesn't represent an absolute address in the device's memory; instead, it's relative to the memory space reserved for that record. The B component has the same 32-bit range, but starts at position #4.

This is a generalized view of the syntax:

These are the elements we see above:

- Component_Name: name of the component (from the record type declaration);
- Start_Position: start position in storage units of the memory space reserved for that component;
- First_Bit: first bit (in the start position) of the component;
- Last_Bit: last bit of the component.

Note that the last bit of a component might be in a different storage unit. Since the **Integer** type has a larger width (32 bits) than the storage unit (8 bits), components of that type span over multiple storage units. Therefore, in our example, the first bit of component A is at position #0, while the last bit is at position #3.

Also note that the last eight bits of component A are bits #24 ... #31. If we think in terms of storage units, this corresponds to bits #0 ... #7 of position #3. However, when specifying the last bit in Ada, we always use the First_Bit value as a reference, not the position where those bits might end up. Therefore, we write **range** 0 ... 31, well knowing that those 32 bits span over four storage units (positions #0 ... #3).

1 In the Ada Reference Manual

13.5.1 Record Representation Clauses⁵¹

2.7.1 Storage Place Attributes

We can retrieve information about the start position, and the first and last bits of a component by using the storage place attributes:

- Position, which retrieves the start position of a component;
- First_Bit, which retrieves the first bit of a component;
- Last_Bit, which retrieves the last bit of a component.

Note, however, that these attributes can only be used with actual records, and not with record types.

We can revisit the previous example and verify how the compiler represents the R type in memory:

```
Listing 26: p.ads
```

```
package P is
package P is
type R is record
A : Integer;
B : Integer;
end record;
end P;
```

```
Listing 27: show storage.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
   with System;
2
3
   with P;
                       use P:
4
5
   procedure Show Storage is
6
      R1 : R;
7
   begin
8
                                          n
      Put Line ("R'Size:
9
                 & R'Size'Image);
10
                                          n
      Put_Line ("R'Object_Size:
11
                 & R'Object_Size'Image);
12
      New Line;
13
14
      Put Line ("System.Storage Unit: "
15
                 & System.Storage Unit'Image);
16
      New Line;
17
18
      Put_Line ("R1.A'Position : "
19
                 & R1.A'Position'Image);
20
      Put Line ("R1.A'First Bit : "
21
                 & R1.A'First Bit'Image);
22
      Put Line ("R1.A'Last Bit : "
23
                 & R1.A'Last Bit'Image);
24
      New Line;
25
26
      Put Line ("R1.B'Position : "
27
                 & R1.B'Position'Image);
28
      Put Line ("R1.B'First Bit : "
29
                 & R1.B'First_Bit'Image);
30
      Put Line ("R1.B'Last Bit : "
31
                 & R1.B'Last Bit'Image);
32
   end Show Storage;
33
```

⁵¹ http://www.ada-auth.org/standards/22rm/html/RM-13-5-1.html

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_

→Storage_Clauses.Storage_Place_Attributes

MD5: 05a402585ce71eb47cf972e68c02835e
```

Build output

show_storage.adb:7:04: warning: variable "R1" is read but never assigned [-gnatwv]

Runtime output

R'Size:64R'Object_Size:64System.Storage_Unit:8R1.A'Position :0R1.A'First_Bit :0R1.A'Last_Bit :31R1.B'Position :4R1.B'First_Bit :0R1.B'Last_Bit :31

First of all, we see that the size of the R type is 64 bits, which can be explained by those two 32-bit integer components. Then, we see that components A and B start at positions #0 and #4, and each one makes use of bits in the range from #0 to #31. This matches the graph we've seen above.

1 In the Ada Reference Manual

• 13.5.2 Storage Place Attributes⁵²

2.7.2 Using Representation Clauses

We can use representation clauses to change the way the compiler handles memory for a record type. For example, let's say we want to have an empty storage unit between components A and B. We can use a representation clause where we specify that component B starts at position #5 instead of #4, leaving an empty byte after component A and before component B:

position	0	1	2	3	4	5	6	7	8
bits	#07	#8#15	#16 #23	#24 #31		#07	#8#15	#16#23	#24 #31
component	А							В	

This is the code that implements that:

Listing 28: p.ads

```
1 package P is
2
3 type R is record
4 A : Integer;
```

(continues on next page)

⁵² http://www.ada-auth.org/standards/22rm/html/RM-13-5-2.html

```
B : Integer;
5
       end record;
6
7
       for R use record
8
          A at 0 range 0 .. 31;
9
          B at 5 range 0 .. 31;
10
       end record;
11
12
   end P;
13
```

Listing 29: show_empty_byte.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with P;
                       use P;
3
4
   procedure Show Empty Byte is
5
   begin
6
      Put Line ("R'Size:
7
                 & R'Size'Image);
8
      Put_Line ("R'Object_Size: "
9
                 & R'Object_Size'Image);
10
   end Show_Empty_Byte;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_

⇔Storage_Clauses.Rep_Clauses_Empty_Byte

MD5: c616e534e95a06f2e8b3052a3e8a9aab
```

Runtime output

R'Size: 72 R'Object_Size: 96

When running the application above, we see that, due to the extra byte in the record representation, the sizes increase. On a typical 64-bit PC, R'Size is now 76 bits, which reflects the additional eight bits that we introduced between components A and B. Depending on the target architecture, you may also see that R'Object_Size is now 96 bits, which is the size the compiler selects as the most appropriate for this record type. As we've mentioned in the previous section, we can use aspects to request a specific size to the compiler. In this case, we could use the Object_Size aspect:

Listing 30: p.ads

```
package P is
1
2
       type R is record
3
          A : Integer;
4
          B : Integer;
5
      end record
6
         with Object Size => 72;
7
8
       for R use record
9
          A at 0 range 0 .. 31;
10
          B at 5 range 0 .. 31;
11
      end record;
12
13
   end P;
14
```

Listing 31: show_empty_byte.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with P;
                      use P;
3
4
   procedure Show Empty Byte is
5
   begin
6
      Put_Line ("R'Size:
7
                 & R'Size'Image);
8
      Put_Line ("R'Object_Size: "
9
                 & R'Object_Size'Image);
10
   end Show_Empty_Byte;
11
```

Code block metadata

Runtime output

R'Size: 72 R'Object_Size: 72

If the code compiles, R'Size and R'Object_Size should now have the same value.

2.7.3 Derived Types And Representation Clauses

In some cases, you might want to modify the memory representation of a record without impacting existing code. For example, you might want to use a record type that was declared in a package that you're not allowed to change. Also, you would like to modify its memory representation in your application. A nice strategy is to derive a type and use a representation clause for the derived type.

We can apply this strategy on our previous example. Let's say we would like to use record type R from package P in our application, but we're not allowed to modify package P — or the record type, for that matter. In this case, we could simply derive R as R_New and use a representation clause for R_New. This is exactly what we do in the specification of the child package P.Rep:

Listing 32: p.ads

```
1 package P is
2
3 type R is record
4 A : Integer;
5 B : Integer;
6 end record;
7
8 end P;
```

Listing 33: p-rep.ads

```
1 package P.Rep is
2
3 type R_New is new R
4 with Object_Size => 72;
5
6 for R New use record
```

(continues on next page)

```
7 A at 0 range 0 .. 31;
8 B at 5 range 0 .. 31;
9 end record;
10
11 end P.Rep;
```

Listing 34: show_empty_byte.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with P;
                       use P;
3
   with P.Rep;
                      use P.Rep;
4
5
   procedure Show_Empty_Byte is
6
   begin
7
                                   п
      Put Line ("R'Size:
8
                 & R'Size'Image);
9
      Put_Line ("R'Object_Size: "
10
                 & R'Object_Size'Image);
11
12
      Put_Line ("R_New'Size:
13
                 & R_New'Size'Image);
14
      Put_Line ("R_New'Object_Size: "
15
                 & R_New'Object_Size'Image);
16
   end Show Empty Byte;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_ →Storage_Clauses.Derived_Rep_Clauses_Empty_Byte MD5: 3ale0837f8bd8250f20fc7b274b869d5

Runtime output

R'Size: 64 R'Object_Size: 64 R_New'Size: 72 R_New'Object_Size: 72

When running this example, we see that the R type retains the memory representation selected by the compiler for the target architecture, while the R_New has the memory representation that we specified.

2.7.4 Representation on Bit Level

A very common application of representation clauses is to specify individual bits of a record. This is particularly useful, for example, when mapping registers or implementing protocols.

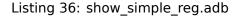
Let's consider the following fictitious register as an example:

bit	0	1	2	3	4	5	6	7
component	C .	5	(rese	rved)	Error		V1	

Here, S is the current status, Error is a flag, and V1 contains a value. Due to the fact that we can use representation clauses to describe individual bits of a register as records, the implementation becomes as simple as this:

```
Listing 35: p.ads
```

```
package P is
1
2
     type Status is (Ready, Waiting,
3
                       Processing, Done);
4
     type UInt 3 is range 0 ... 2 ** 3 - 1;
5
6
      type Simple_Reg is record
7
               : Status;
          S
8
         Error : Boolean;
9
         V1 : UInt_3;
10
      end record;
11
12
      for Simple_Reg use record
13
               at 0 range 0 .. 1;
14
          S
          -- Bit #2 and 3: reserved!
15
          Error at 0 range 4 .. 4;
16
         V1
                at 0 range 5 ... 7;
17
      end record;
18
19
   end P;
20
```



```
with Ada.Text IO; use Ada.Text IO;
1
2
3
   with P;
                      use P;
4
   procedure Show Simple Reg is
5
   begin
6
      Put_Line ("Simple Reg'Size:
7
                 & Simple Reg'Size'Image);
8
      Put_Line ("Simple_Reg'Object_Size: "
9
                 & Simple Reg'Object Size'Image);
10
   end Show Simple Reg;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_

→Storage_Clauses.Rep_Clauses_Simple_Reg

MD5: cbac444336572460062f922767c226a5
```

Runtime output

Simple_Reg'Size: 8 Simple_Reg'Object_Size: 8

As we can see in the declaration of the Simple_Reg type, each component represents a field from our register, and it has a fixed location (which matches the register representation we see in the graph above). Any operation on the register is as simple as accessing the record component. For example:

Listing	37:	show	simple	_reg.adb

1	with Ac	da.Text IO:	use Ada.Text IO;	
2		_ ,	,	
3	with P;	;	use P;	
				(continues on next page)

```
4
   procedure Show_Simple_Reg is
5
      Default : constant Simple Reg :=
6
                   (S => Ready,
7
                    Error => False,
8
                    V1
                          => 0);
9
10
      R : Simple_Reg := Default;
11
   begin
12
      Put_Line ("R.S: " & R.S'Image);
13
14
      R.V1 := 4;
15
16
      Put_Line ("R.V1: " & R.V1'Image);
17
   end Show_Simple_Reg;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Record_Representation_

→Storage_Clauses.Rep_Clauses_Simple_Reg

MD5: e442396e43d6609c1c837165bbc21641
```

Runtime output

R.S: READY R.V1: 4

As we can see in the example, to retrieve the current status of the register, we just have to write R.S. To update the V1 field of the register with the value 4, we just have to write R.V1 := 4. No extra code — such as bit-masking or bit-shifting — is needed here.

In other languages

Some programming languages require that developers use complicated, error-prone approaches — which may include manually bit-shifting and bit-masking variables — to retrieve information from or store information to individual bits or registers. In Ada, however, this is efficiently handled by the compiler, so that developers only need to correctly describe the register mapping using representation clauses.

2.8 Changing Data Representation

\rm Note

This section was originally written by Robert Dewar and published as Gem #27: Changing Data Representation⁵³ and Gem #28⁵⁴.

A powerful feature of Ada is the ability to specify the exact data layout. This is particularly important when you have an external device or program that requires a very specific format. Some examples are:

⁵³ https://www.adacore.com/gems/gem-27

⁵⁴ https://www.adacore.com/gems/gem-28

```
Listing 38: communication.ads
```

```
package Communication is
1
2
      type Com_Packet is record
3
         Key : Boolean;
4
          Id : Character;
5
          Val : Integer range 100 .. 227;
6
      end record;
7
8
      for Com Packet use record
9
          Key at 0 range 0 .. 0;
10
          Id at 0 range 1 .. 8;
11
         Val at 0 range 9 .. 15;
12
      end record;
13
14
   end Communication:
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

→Representation.Com_Packet

MD5: cbd7f5547c5b0458853ac21d03aa41f8
```

Build output

which lays out the fields of a record, and in the case of Val, forces a biased representation in which all zero bits represents 100. Another example is:

Listing 39: array representation.ads

```
package Array_Representation is
type Val is (A, B, C, D, E, F, G, H);
type Arr is array (1 .. 16) of Val
with Component_Size => 3;
end Array Representation;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

→Representation.Array_Rep

MD5: 7eb17fc2cd415acb7c53a363fa336807
```

which forces the components to take only 3 bits, crossing byte boundaries as needed. A final example is:

Listing 40: enumeration_representation.ads

```
package Enumeration_Representation is
1
2
      type Status is (Off, On, Unknown);
3
      for Status use (Off
                             => 2#001#,
4
                       0n
                              => 2#010#,
5
                      Unknown => 2#100#);
6
7
8
  end Enumeration_Representation;
```

Code block metadata

which allows specified values for an enumeration type, instead of the efficient default values of 0, 1, 2.

In all these cases, we might use these representation clauses to match external specifications, which can be very useful. The disadvantage of such layouts is that they are inefficient, and accessing individual components, or, in the case of the enumeration type, looping through the values can increase space and time requirements for the program code.

One approach that is often effective is to read or write the data in question in this specified form, but internally in the program represent the data in the normal default layout, allowing efficient access, and do all internal computations with this more efficient form.

To follow this approach, you will need to convert between the efficient format and the specified format. Ada provides a very convenient method for doing this, as described in RM 13.6 "Change of Representation"⁵⁵.

The idea is to use type derivation, where one type has the specified format and the other has the normal default format. For instance for the array case above, we would write:

Listing 41: array_representation.ads

```
package Array_Representation is
    type Val is (A, B, C, D, E, F, G, H);
    type Arr is array (1 .. 16) of Val;
    type External_Arr is new Arr
    with Component_Size => 3;
    end Array_Representation;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

⇔Representation.Array_Rep

MD5: d4e90f6ef8ff81771980771356eab235
```

Now we read and write the data using the External_Arr type. When we want to convert to the efficient form, Arr, we simply use a type conversion.

Listing 42: using array for io.adb

```
with Arrav Representation:
1
   use Array Representation;
2
3
   procedure Using Array For IO is
4
      Input_Data : External Arr;
5
      Work Data
                 : Arr;
6
      Output_Data : External_Arr;
7
   begin
8
          (read data into Input Data)
9
10
      -- Now convert to internal form
11
      Work Data := Arr (Input Data);
12
13
```

(continues on next page)

⁵⁵ http://www.ada-auth.org/standards/22rm/html/RM-13-6.html

```
14 -- (computations using efficient
15 -- Work_Data form)
16
17 -- Convert back to external form
18 Output_Data := External_Arr (Work_Data);
19
20 end Using_Array_For_IO;
```

Code block metadata

Build output

```
using_array_for_io.adb:5:04: warning: variable "Input_Data" is read but never

_assigned [-gnatwv]
```

Using this approach, the quite complex task of copying all the data of the array from one form to another, with all the necessary masking and shift operations, is completely automatic.

Similar code can be used in the record and enumeration type cases. It is even possible to specify two different representations for the two types, and convert from one form to the other, as in:

```
Listing 43: enumeration_representation.ads
```

```
package Enumeration Representation is
1
2
       type Status In is (Off, On, Unknown);
3
       type Status Out is new Status In;
4
5
      for Status In use (Off
                                    => 2#001#,
6
                           0n
                                    => 2#010#.
7
                           Unknown => 2#100#);
8
       for Status Out use (Off
                                     => 103.
9
                             0n
                                     => 1045
10
                            Unknown \Rightarrow 7700);
11
12
   end Enumeration Representation;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

←Representation.Enum_Rep

MD5: f78c3718280f9265ff54270c5834b458
```

There are two restrictions that must be kept in mind when using this feature. First, you have to use a derived type. You can't put representation clauses on subtypes, which means that the conversion must always be explicit. Second, there is a rule RM 13.1^{56} (10) that restricts the placement of interesting representation clauses:

10 For an untagged derived type, no type-related representation items are allowed if the parent type is a by-reference type, or has any user-defined primitive subprograms.

All the representation clauses that are interesting from the point of view of change of representation are "type related", so for example, the following sequence would be illegal:

⁵⁶ http://www.ada-auth.org/standards/22rm/html/RM-13-1.html

Listing 44: array_representation.ads

```
package Array Representation is
1
2
      type Val is (A, B, C, D, E, F, G, H);
3
      type Arr is array (1 .. 16) of Val;
4
5
      procedure Rearrange (Arg : in out Arr);
6
7
      type External Arr is new Arr
8
        with Component Size => 3;
9
10
   end Array Representation;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

→Representation.Array_Rep_2

MD5: 70201932d40e3fb356bc1d8ab188f2df
```

Build output

```
array_representation.ads:9:11: error: representation item not permitted before Ada

△2022

array_representation.ads:9:11: error: parent type "Arr" has primitive operations

gprbuild: *** compilation phase failed
```

Why these restrictions? Well, the answer is a little complex, and has to do with efficiency considerations, which we will address below.

2.8.1 Restrictions

In the previous subsection, we discussed the use of derived types and representation clauses to achieve automatic change of representation. More accurately, this feature is not completely automatic, since it requires you to write an explicit conversion. In fact there is a principle behind the design here which says that a change of representation should never occur implicitly behind the back of the programmer without such an explicit request by means of a type conversion.

The reason for that is that the change of representation operation can be very expensive, since in general it can require component by component copying, changing the representation on each component.

Let's have a look at the -gnatG expanded code to see what is hidden under the covers here. For example, the conversion Arr (Input_Data) from the previous example generates the following expanded code:

```
B26b : declare
[subtype p__TarrD1 is integer range 1 .. 16]
R25b : p__TarrD1 := 1;
begin
for L24b in 1 .. 16 loop
[subtype p__arr___XP3 is
system__unsigned_types__long_long_unsigned range 0 ..
16#FFFF_FFFFFF#]
work_data := p__arr___XP3!((work_data and not shift_left!(
16#7#, 3 * (integer(L24b - 1)))) or shift_left!(p__arr___XP3!
(input_data (R25b)), 3 * (integer(L24b - 1))));
R25b := p__TarrD1'succ(R25b);
end loop;
end B26b;
```

That's pretty horrible! In fact, we could have simplified it for this section, but we have left it in its original form, so that you can see why it is nice to let the compiler generate all this stuff so you don't have to worry about it yourself.

Given that the conversion can be pretty inefficient, you don't want to convert backwards and forwards more than you have to, and the whole approach is only worthwhile if we'll be doing extensive computations involving the value.

The expense of the conversion explains two aspects of this feature that are not obvious. First, why do we require derived types instead of just allowing subtypes to have different representations, avoiding the need for an explicit conversion?

The answer is precisely that the conversions are expensive, and you don't want them happening behind your back. So if you write the explicit conversion, you get all the gobbledygook listed above, but you can be sure that this never happens unless you explicitly ask for it.

This also explains the restriction we mentioned in previous subsection from RM 13.1⁵⁷ (10):

10 For an untagged derived type, no type-related representation items are allowed if the parent type is a by-reference type, or has any user-defined primitive subprograms.

It turns out this restriction is all about avoiding implicit changes of representation. Let's have a look at how type derivation works when there are primitive subprograms defined at the point of derivation. Consider this example:

Listing 45: my_ints.ads

```
package My Ints is
1
2
      type My_Int_1 is range 1 .. 10;
3
4
      function Odd (Arg : My_Int_1)
5
                      return Boolean;
6
7
      type My Int 2 is new My Int 1;
8
9
   end My_Ints;
10
```

Listing 46: my ints.adb

```
1 package body My_Ints is
2
3 function Odd (Arg : My_Int_1)
4 return Boolean is
5 (True);
6 -- Dummy implementation!
7
8 end My_Ints;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

□Representation.My_Int

MD5: a29401698307998288f02b349d04d1d2
```

Now when we do the type derivation, we inherit the function Odd for My_Int_2. But where does this function come from? We haven't written it explicitly, so the compiler somehow materializes this new implicit function. How does it do that?

We might think that a complete new function is created including a body in which My_Int_2

⁵⁷ http://www.ada-auth.org/standards/22rm/html/RM-13-1.html

replaces My_Int_1, but that would be impractical and expensive. The actual mechanism avoids the need to do this by use of implicit type conversions. Suppose after the above declarations, we write:

```
Listing 47: using_my_int.adb
```

```
with My_Ints; use My_Ints;
1
2
   procedure Using My Int is
3
      Var : My_Int_2;
4
   begin
5
6
      if Odd (Var) then
7
                ^ Calling Odd function
8
          - -
          - -
                 for My_Int_2 type.
9
          null;
10
      end if;
11
12
   end Using_My_Int;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

⊲Representation.My_Int

MD5: f68272d55e68687b7102885313c7831b
```

Build output

using_my_int.adb:4:04: warning: variable "Var" is read but never assigned [-gnatwv]

The compiler translates this as:

Listing 48: using_my_int.adb

```
with My_Ints; use My_Ints;
1
2
   procedure Using_My_Int is
3
      Var : My_Int_2;
4
   begin
5
6
      if Odd (My_Int_1 (Var)) then
7
               ^ Converting My_Int_2 to
8
          - -
                 My_Int_1 type before
          - -
9
                 calling Odd function.
          - -
10
          null:
11
      end if;
12
13
  end Using_My_Int;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Changing_Data_

⇔Representation.My_Int

MD5: b3d0053c61412a2b985cd580b645e048
```

Build output

using_my_int.adb:4:04: warning: variable "Var" is read but never assigned [-gnatwv]

This implicit conversion is a nice trick, it means that we can get the effect of inheriting a new operation without actually having to create it. Furthermore, in a case like this, the type conversion generates no code, since My_Int_1 and My_Int_2 have the same representation.

But the whole point is that they might not have the same representation if one of them had a representation clause that made the representations different, and in this case the implicit conversion inserted by the compiler could be expensive, perhaps generating the junk we quoted above for the Arr case. Since we never want that to happen implicitly, there is a rule to prevent it.

The business of forbidding by-reference types (which includes all tagged types) is also driven by this consideration. If the representations are the same, it is fine to pass by reference, even in the presence of the conversion, but if there was a change of representation, it would force a copy, which would violate the by-reference requirement.

So to summarize this section, on the one hand Ada gives you a very convenient way to trigger these complex conversions between different representations. On the other hand, Ada guarantees that you never get these potentially expensive conversions happening unless you explicitly ask for them.

2.9 Valid Attribute

When receiving data from external sources, we're subjected to problems such as transmission errors. If not handled properly, erroneous data can lead to major issues in an application.

One of those issues originates from the fact that transmission errors might lead to invalid information stored in memory. When proper checks are active, using invalid information is detected at runtime and an exception is raised at this point, which might then be handled by the application.

Instead of relying on exception handling, however, we could instead ensure that the information we're about to use is valid. We can do this by using the Valid attribute. For example, if we have a variable Var, we can verify that the value stored in Var is valid by writing Var'Valid, which returns a Boolean value. Therefore, if the value of Var isn't valid, Var'Valid returns False, so we can have code that handles this situation before we actually make use of Var. In other words, instead of handling a potential exception in other parts of the application, we can proactively verify that input information is correct and avoid that an exception is raised.

In the next example, we show an application that

- generates a file containing mock-up data, and then
- reads information from this file as state values.

The mock-up data includes valid and invalid states.

```
Listing 49: create test file.ads
```

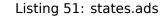
```
procedure Create_Test_File (File_Name : String);
```

Listing 50: create_test_file.adb

```
with Ada.Sequential IO;
1
2
   procedure Create_Test_File (File_Name : String)
3
   is
4
      package Integer Sequential IO is new
5
        Ada.Sequential IO (Integer);
6
      use Integer_Sequential_I0;
7
8
      F : File_Type;
9
   begin
10
      Create (F, Out_File, File_Name);
11
```

(continues on next page)

```
Write (F,
                   1);
12
       Write (F,
                   2);
13
       Write (F,
                   4);
14
       Write (F,
                   3);
15
       Write (F,
                  2);
16
       Write (F,
                   10);
17
       Close (F);
18
   end Create_Test_File;
19
```



```
with Ada.Sequential_I0;
1
2
   package States is
3
4
      type State is (Off, On, Waiting)
5
         with Size => Integer'Size;
6
7
      for State use (Off
                                => 1,
8
                                => 2,
                       0n
9
                       Waiting => 4);
10
11
      package State_Sequential_I0 is new
12
         Ada.Sequential_IO (State);
13
14
      procedure Read_Display_States
15
         (File_Name : String);
16
17
18
   end States;
```

Listing 52: states.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
3
   package body States is
4
      procedure Read_Display_States
5
         (File_Name : String)
6
      is
7
          use State_Sequential_I0;
8
9
          F : State_Sequential_I0.File_Type;
10
          S : State;
11
12
          procedure Display_State (S : State) is
13
          begin
14
             - -
                 Before displaying the value,
15
                 check whether it's valid or not.
16
             - -
             if S'Valid then
17
                Put_Line (S'Image);
18
             else
19
                Put_Line ("Invalid value detected!");
20
             end if;
21
          end Display_State;
22
23
24
      begin
          Open (F, In_File, File_Name);
25
26
          while not End_Of_File (F) loop
27
             Read (F, S);
28
```

(continues on next page)

```
29 Display_State (S);
30 end loop;
31
32 Close (F);
33 end Read_Display_States;
34
35 end States;
```

Listing 53: show_states_from_file.adb

```
with States;
                           use States;
1
  with Create_Test_File;
2
3
   procedure Show_States_From File is
4
      File Name : constant String := "data.bin";
5
  begin
6
      Create Test File (File Name);
7
      Read Display States (File Name);
8
  end Show_States_From_File;
```

Code block metadata

Runtime output

OFF ON WAITING Invalid value detected! ON Invalid value detected!

Let's start our discussion on this example with the States package, which contains the declaration of the State type. This type is a simple enumeration containing three states: Off, On and Waiting. We're assigning specific integer values for this type by declaring an enumeration representation clause. Note that we're using the Size aspect to request that objects of this type have the same size as the **Integer** type. This becomes important later on when parsing data from the file.

In the Create_Test_File procedure, we create a file containing integer values, which is parsed later by the Read_Display_States procedure. The Create_Test_File procedure doesn't contain any reference to the State type, so we're not constrained to just writing information that is valid for this type. On the contrary, this procedure makes use of the **Integer** type, so we can write any integer value to the file. We use this strategy to write both valid and invalid values of State to the file. This allows us to simulate an environment where transmission errors occur.

We call the Read_Display_States procedure to read information from the file and display each state stored in the file. In the main loop of this procedure, we call Read to read a state from the file and store it in the S variable. We then call the nested Display_State procedure to display the actual state stored in S. The most important line of code in the Display_State procedure is the one that uses the Valid attribute:

if S'Valid then

In this line, we're verifying that the S variable contains a valid state before displaying the actual information from S. If the value stored in S isn't valid, we can handle the issue accordingly. In this case, we're simply displaying a message indicating that an invalid value was

detected. If we didn't have this check, the Constraint_Error exception would be raised when trying to use invalid data stored in S — this would happen, for example, after reading the integer value 3 from the input file.

In summary, using the Valid attribute is a good strategy we can employ when we know that information stored in memory might be corrupted.

In the Ada Reference Manual

• 13.9.2 The Valid Attribute⁵⁸

2.10 Unchecked Union

We've introduced variant records back in the Introduction to Ada course⁵⁹. In simple terms, a variant record is a record with discriminants that allows for changing its structure. Basically, it's a record containing a **case**. (We talk again about *variant records* (page 227) in another chapter.)

The State_Or_Integer declaration in the States package below is an example of a variant record:

Listing 54: states.ads

```
package States is
1
2
      type State is (Off, On, Waiting)
3
         with Size => Integer'Size;
4
5
      for State use (Off
                               => 1,
6
                               => 2,
                       0n
7
                       Waiting => 4);
8
9
      type State_Or_Integer (Use_Enum : Boolean) is
10
       record
11
          case Use Enum is
12
             when False => I : Integer;
13
             when True => S : State;
14
          end case;
15
      end record;
16
17
      procedure Display_State_Value
18
         (V : State_Or_Integer);
19
20
   end States;
21
```

Listing 55: states.adb

```
with Ada.Text_I0; use Ada.Text_I0;
package body States is
procedure Display_State_Value
(V : State_Or_Integer)
is
begin
```

(continues on next page)

⁵⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/more_about_records.html# intro-ada-variant-records

⁵⁸ http://www.ada-auth.org/standards/22rm/html/RM-13-9-2.html

```
9 Put_Line ("State: " & V.S'Image);
10 Put_Line ("Value: " & V.I'Image);
11 end Display_State_Value;
12
13 end States;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Unchecked_Union.State_ Gor_Integer MD5: fa72f52a4396a2e66931ff6932c567fc

As mentioned in the previous course, if you try to access a component that is not valid for your record, a Constraint_Error exception is raised. For example, in the implementation of the Display_State_Value procedure, we're trying to retrieve the value of the integer component (I) of the V record. When calling this procedure, the Constraint_Error exception is raised as expected because Use_Enum is set to **True**, so that the I component is invalid — only the S component is valid in this case.

Listing 56: show_variant_rec_error.adb

```
with States; use States;
procedure Show_Variant_Rec_Error is
V : State_Or_Integer (Use_Enum => True);
begin
V.S := On;
Display_State_Value (V);
end Show_Variant_Rec_Error;
```

Code block metadata

Runtime output

State: ON

raised CONSTRAINT_ERROR : states.adb:10 discriminant check failed

In addition to not being able to read the value of a component that isn't valid, assigning a value to a component that isn't valid also raises an exception at runtime. In this example, we cannot assign to V.I:

Listing 57: show_variant_rec_error.adb

```
with States; use States;
1
2
   procedure Show_Variant_Rec_Error is
3
      V : State_Or_Integer (Use_Enum => True);
4
  begin
5
      V.I := 4;
6
      -- Error: V.I cannot be accessed because
7
                 Use Enum is set to True.
8
  end Show_Variant_Rec_Error;
9
```

Code block metadata

Build output

Runtime output

raised CONSTRAINT_ERROR : show_variant_rec_error.adb:6 discriminant check failed

We may circumvent this limitation by using the Unchecked_Union aspect. For example, we can derive a new type from State_Or_Integer and use this aspect in its declaration. We do this in the declaration of the Unchecked_State_Or_Integer type below.

Listing	58:	states.ads

```
package States is
1
2
       type State is (Off, On, Waiting)
3
         with Size => Integer'Size;
4
5
       for State use (Off
                                => 1,
6
                                => 2,
                       0n
7
                       Waiting => 4);
8
9
       type State_Or_Integer (Use_Enum : Boolean) is
10
       record
11
          case Use Enum is
12
             when False => I : Integer;
13
             when True => S : State;
14
          end case;
15
       end record;
16
17
       type Unchecked_State_Or_Integer
18
         (Use_Enum : Boolean) is new
19
           State_Or_Integer (Use_Enum)
20
             with Unchecked_Union;
21
22
       procedure Display State Value
23
         (V : Unchecked_State_Or_Integer);
24
25
   end States;
26
```

Listing 59: states.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body States is
3
4
      procedure Display_State_Value
5
         (V : Unchecked_State_Or_Integer)
6
      is
7
      begin
8
         Put_Line ("State: " & V.S'Image);
9
         Put Line ("Value: " & V.I'Image);
10
```

(continues on next page)

ii end Display_State_Value;

```
12
13
```

end States;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Unchecked_Union.

→Unchecked_State_Or_Integer

MD5: e97271a24aab23d2db450308401667ac
```

Because we now use the Unchecked_State_Or_Integer type for the input parameter of the Display_State_Value procedure, no exception is raised at runtime, as both components are now accessible. For example:

Listing 60: show_unchecked_union.adb

```
with States; use States;
1
2
  procedure Show Unchecked Union is
3
      V : State Or Integer (Use Enum => True);
4
  begin
5
      V.S := 0n;
6
      Display_State_Value
7
        (Unchecked State Or Integer (V));
8
  end Show Unchecked Union;
9
```

Code block metadata

Runtime output

State: ON Value: 2

Note that, in the call to the Display_State_Value procedure, we first need to convert the V argument from the State_Or_Integer to the Unchecked_State_Or_Integer type.

Also, we can assign to any of the components of a record that has the Unchecked_Union aspect. In our example, we can now assign to both the S and the I components of the V record:

Listing 61: show_unchecked_union.adb

```
with States; use States;
1
2
   procedure Show Unchecked Union is
3
      V : Unchecked_State_Or_Integer
4
             (Use_Enum => True);
5
   begin
6
      V := (Use Enum => True, S => On);
7
      Display_State_Value (V);
8
9
      V := (Use Enum => False, I => 4);
10
      Display State Value (V);
11
   end Show Unchecked Union;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Unchecked_Union. →Unchecked_State_Or_Integer MD5: bb472e91c5e7b7e63d6246dbcf5226a0

Runtime output

State: ON Value: 2 State: WAITING Value: 4

In the example above, we're use an aggregate in the assignments to V. By doing so, we avoid that Use_Enum is set to the *wrong* component. For example:

Listing 62: show unchecked union.adb

```
with States; use States;
1
2
   procedure Show Unchecked Union is
3
      V : Unchecked State Or Integer
4
             (Use_Enum => True);
5
   begin
6
      V.S := 0n;
7
      Display_State_Value (V);
8
9
      V.I := 4:
10
      -- Error: cannot directly assign to V.I,
11
                  as Use Enum is set to True.
12
13
      Display State Value (V);
14
   end Show Unchecked Union;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Unchecked_Union.

GUnchecked_State_Or_Integer

MD5: 74ac11a3effdafd3959fface295a86da
```

Build output

Runtime output

State: ON Value: 2

raised CONSTRAINT_ERROR : show_unchecked_union.adb:10 discriminant check failed

Here, even though the record has the Unchecked_Union attribute, we cannot directly assign to the I component because Use_Enum is set to **True**, so only the S is accessible. We can, however, read its value, as we do in the Display_State_Value procedure.

Be aware that, due to the fact the union is not checked, we might write invalid data to the record. In the example below, we initialize the I component with 3, which is a valid integer value, but results in an invalid value for the S component, as the value 3 cannot be mapped to the representation of the State type.

```
Listing 63: show_unchecked_union.adb
```

```
with States; use States;
1
2
   procedure Show_Unchecked_Union is
3
      V : Unchecked_State_Or_Integer
4
            (Use_Enum => True);
5
   begin
6
      V := (Use Enum => False, I => 3);
7
      Display_State_Value (V);
8
  end Show Unchecked Union;
9
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Unchecked_Union.

□Unchecked_State_Or_Integer

MD5: f63e64df137cfc3c29e41f784306f0e4
```

Runtime output

raised CONSTRAINT_ERROR : states.adb:9 invalid data

To mitigate this problem, we could use the Valid attribute — discussed in the previous section — for the S component before trying to use its value in the implementation of the Display_State_Value procedure:

```
Listing 64: states.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body States is
3
4
      procedure Display State Value
5
         (V : Unchecked State Or Integer)
6
      is
7
      begin
8
         if V.S'Valid then
9
             Put_Line ("State: " & V.S'Image);
10
         else
11
             Put_Line ("State: <invalid>");
12
         end if;
13
         Put_Line ("Value: " & V.I'Image);
14
      end Display_State_Value;
15
16
   end States;
17
```

Listing 65: show_unchecked_union.adb

```
with States; use States;
1
2
   procedure Show_Unchecked_Union is
3
      V : Unchecked_State_Or_Integer
4
            (Use Enum => True);
5
  beain
6
      V := (Use_Enum => False, I => 3);
7
      Display State Value (V);
8
  end Show_Unchecked_Union;
9
```

However, in general, you should avoid using the Unchecked_Union aspect due to the potential issues you might introduce into your application. In the majority of the cases, you don't need it at all — except for special cases such as when interfacing with C code that makes use of union types or solving very specific problems when doing low-level programming.

- **1** In the Ada Reference Manual
 - B.3.3 Unchecked Union Types⁶⁰

2.11 Addresses

In other languages, such as C, the concept of pointers and addresses plays a prominent role. (In fact, in C, many optimizations rely on the usage of pointer arithmetic.) The concept of addresses does exist in Ada, but it's mainly reserved for very specific applications, mostly related to low-level programming. In general, other approaches — such as using access types — are more than sufficient. (We discuss *access types* (page 593) in another chapter. Also, later on in that chapter, we discuss the *relation between access types and addresses* (page 706).) In this section, we discuss some details about using addresses in Ada.

We make use of the Address type, which is defined in the System package, to handle addresses. In contrast to other programming languages (such as C or C++), an address in Ada isn't an integer value: its definition depends on the compiler implementation, and it's actually driven directly by the hardware. For now, let's consider it to usually be a private type — this can be seen as an attempt to achieve application code portability, given the variations in hardware that result in different definitions of what an address actually is.

The Address type has support for *address comparison* (page 126) and *address arithmetic* (page 128) (also known as *pointer arithmetic* in C). We discuss these topics later in this section. First, let's talk about the Address attribute and the Address aspect.

1 In the Ada Reference Manual

```
• 13.7 The Package System<sup>61</sup>
```

2.11.1 Address attribute

The **Address** attribute allows us to get the address of an object. For example:

Listing 66: use address.adb

```
with System; use System;
procedure Use_Address is
    I : aliased Integer := 5;
    A : Address;
    begin
    A := I'Address;
    end Use_Address;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Address_

⇔Attribute

MD5: 1ee71b7cd3ed278647eb72f383da877f
```

Here, we're assigning the address of the I object to the A address.

⁶⁰ http://www.ada-auth.org/standards/22rm/html/RM-B-3-3.html

⁶¹ http://www.ada-auth.org/standards/22rm/html/RM-13-7.html

1 In the GNAT toolchain

GNAT offers a very useful extension to the System package to retrieve a string for an address: System.Address_Image. This is the function profile:

```
function System.Address_Image
  (A : System.Address) return String;
```

with Ada.Text IO; use Ada.Text IO;

with System.Address Image;

We can use this function to display the address in an user message, for example:

Listing 67: show address attribute.adb

```
1
2
3
4
5
6
7
8
```

9

procedure Show_Address_Attribute is
 I : aliased Integer := 5;
begin
 Put_Line ("Address : "
 & System.Address_Image (I'Address));
end Show Address Attribute;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Show_

⇔Address_Attribute

MD5: 72efddedc57701665594de5ee1939d3d
```

Runtime output

Address : 00007FFEE20218C4

1 In the Ada Reference Manual

- 13.3 Operational and Representation Attributes⁶²
- 13.7 The Package System⁶³

2.11.2 Address aspect

Usually, we let the compiler select the address of an object in memory, or let it use a register to store that object. However, we can specify the address of an object with the **Address** aspect. In this case, the compiler won't select an address automatically, but use the address that we're specifying. For example:

```
Listing 68: show address.adb
```

```
with System; use System;
1
   with System.Address Image;
2
   with Ada.Text_I0; use Ada.Text_I0;
5
   procedure Show Address is
6
7
      I Main
              : aliased Integer;
8
      I Mapped : Integer
9
                    with Address => I_Main'Address;
10
```

(continues on next page)

⁶² http://www.ada-auth.org/standards/22rm/html/RM-13-3.html

⁶³ http://www.ada-auth.org/standards/22rm/html/RM-13-7.html

```
begin
11
      Put_Line ("I_Main'Address
                                      : "
12
                   & System.Address_Image
13
                       (I_Main'Address));
14
      Put_Line ("I_Mapped'Address :
15
                  & System.Address_Image
16
                       (I_Mapped'Address));
17
   end Show Address;
18
```

Code block metadata

Runtime output

I_Main'Address : 00007FFE4E6ACA74 I_Mapped'Address : 00007FFE4E6ACA74

This approach allows us to create an overlay. For example:

Listing 69: simple_overlay.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Simple Overlay is
3
      type State is (Off, State 1, State 2)
4
        with Size => Integer'Size;
5
6
      for State use (Off
                              => 0,
7
                       State 1 => 32,
8
                       State 2 \implies 64;
9
10
      S : State;
11
      I : Integer
12
        with Address => S'Address, Import, Volatile;
13
   begin
14
      S := State 2;
15
      Put Line ("I = " & Integer'Image (I));
16
   end Simple Overlay;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Simple_

Goverlay

MD5: a65057882518824d3ea173d193a7ae67
```

Runtime output

I = 64

Here, I is an overlay of S, as it uses S'Address. With this approach, we can either use the enumeration directly (by using the S object of State type) or its integer representation (by using the I variable).

In the GNAT toolchain

We could call the GNAT-specific System'To_Address attribute when using the Address aspect:

```
Listing 70: shared_var_types.ads
   with System;
1
2
   package Shared_Var_Types is
3
4
   private
5
      R : Integer
6
             with Atomic,
7
8
                  Address =>
                    System'To_Address (16#FFFF00A0#);
9
10
   end Shared_Var_Types;
11
   Code block metadata
   Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Show_
    →Access Address
   MD5: 5c2d8e0a9615084c2a15f896c61adaa6
   In this case, R will refer to the address in memory that we're specifying (16#FFFF00A0#
   in this case).
   As explained in the GNAT Reference Manual<sup>64</sup>, the System'To Address attribute denotes
   a function identical to To Address (from the System.Storage Elements package) ex-
   cept that it is a static attribute. (We talk about the To Address function (page 127)
```

function later on.) Note that we're using the Atomic aspect here, which we discuss *in another chapter* (page 142).

```
In the Ada Reference Manual
```

- 13.3 Operational and Representation Attributes⁶⁵
- 13.7 The Package System⁶⁶
- 13.7.1 The Package System.Storage_Elements⁶⁷

2.11.3 Address comparison

We can compare addresses using the common comparison operators. For example:

Listing 71: show_address.adb

```
with System; use System;
1
   with System.Address_Image;
2
3
   with Ada.Text_IO; use Ada.Text_IO;
4
5
  procedure Show_Address is
6
7
      I, J : Integer;
8
   begin
9
      Put_Line ("I'Address : "
10
                  & System.Address_Image
11
                      (I'Address));
12
```

(continues on next page)

⁶⁴ https://gcc.gnu.org/onlinedocs/gnat_rm/Attribute-To_005fAddress.html

⁶⁵ http://www.ada-auth.org/standards/22rm/html/RM-13-3.html

⁶⁶ http://www.ada-auth.org/standards/22rm/html/RM-13-7.html

⁶⁷ http://www.ada-auth.org/standards/22rm/html/RM-13-7-1.html

```
Put_Line ("J'Address
                                 11 <sup>10</sup>
13
                    & System.Address_Image
14
                        (J'Address));
15
16
       if I'Address = J'Address then
17
          Put_Line ("I'Address = J'Address");
18
       elsif I'Address < J'Address then
19
          Put Line ("I'Address < J'Address");</pre>
20
       else
21
          Put_Line ("I'Address > J'Address");
22
       end if;
23
   end Show_Address;
24
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Address_

⇔Aspect

MD5: 24ddb7d05159f26ef3b2ff6bcc2691e8
```

Runtime output

I'Address : 00007FFC2E3E07FC J'Address : 00007FFC2E3E07F8 I'Address > J'Address

In this example, we compare the address of the I object with the address of the J object using the =, < and > operators.

In the Ada Reference Manual

```
    13.7 The Package System<sup>68</sup>
```

2.11.4 Address to integer conversion

The System.Storage_Elements package offers an integer representation of an address via the Integer_Address type, which is an integer type unrelated to common integer types such as **Integer** and **Long_Integer**. (The actual definition of Integer_Address is compiler-dependent, and it can be a signed or modular integer subtype.)

We can convert between the Address and Integer_Address types by using the To_Address and To_Integer functions. Let's see an example:

```
Listing 72: show_address.adb
```

```
with System;
                      use System;
1
2
   with System.Storage Elements;
3
   use System.Storage_Elements;
4
5
   with System.Address_Image;
6
7
8
   with Ada.Text_I0; use Ada.Text_I0;
9
   procedure Show_Address is
10
      Т
             : Integer;
11
      A1, A2 : Address;
12
           : Integer_Address;
      IΑ
13
```

(continues on next page)

⁶⁸ http://www.ada-auth.org/standards/22rm/html/RM-13-7.html

```
begin
14
      A1 := I'Address;
15
      IA := To_Integer (A1);
16
      A2 := To_Address (IA);
17
18
      Put_Line ("A1 : "
19
                   & System.Address_Image (A1));
20
      Put_Line ("IA :
21
                  & Integer_Address'Image (IA));
22
      Put_Line ("A2 :
23
                  & System.Address_Image (A2));
24
   end Show_Address;
25
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Pointer_ ⇔Arith_Ada MD5: 69e053886fb8e8571d6c94247dc9f30f

Runtime output

A1 : 00007FFF6DCC5BF4 IA : 140735035497460 A2 : 00007FFF6DCC5BF4

Here, we retrieve the address of the I object and store it in the A1 address. Then, we convert A1 to an integer address by calling To_Integer (and store it in IA). Finally, we convert this integer address back to an actual address by calling To_Address.

```
1 In the Ada Reference Manual
```

• 13.7.1 The Package System.Storage_Elements⁶⁹

2.11.5 Address arithmetic

Although Ada supports address arithmetic, which we discuss in this section, it should be reserved for very specific applications such as low-level programming. However, even in situations that require close access to the underlying hardware, using address arithmetic might not be the approach you should consider — make sure to evaluate other options first!

Ada supports address arithmetic via the System.Storage_Elements package, which includes operators such as + and - for addresses. Let's see a code example where we iterate over an array by incrementing an address that *points* to each component in memory:

Listing	73:	show	address	.adb

```
with System;
                      use System;
1
2
   with System.Storage_Elements;
3
   use System.Storage_Elements;
4
5
   with System.Address_Image;
6
7
   with Ada.Text_IO; use Ada.Text_IO;
8
9
  procedure Show_Address is
10
11
```

(continues on next page)

⁶⁹ http://www.ada-auth.org/standards/22rm/html/RM-13-7-1.html

```
Arr : array (1 .. 10) of Integer;
12
          : Address := Arr'Address;
13
       Α
       - -
14
            Initializing address object with
       - -
15
           address of the first component of Arr.
16
       - -
17
       - -
           We could write this as well:
18
       - -
            ____ := Arr (1) 'Address
       - -
19
20
   begin
21
       for I in Arr'Range loop
22
          declare
23
             Curr : Integer
24
                       with Address => A;
25
          begin
26
             Curr := I;
27
             Put_Line ("Curr'Address : "
28
                        & System.Address_Image
29
                             (Curr'Address));
30
          end;
31
32
33
          -- Address arithmetic
34
35
          - -
          A := A + Storage_Offset (Integer'Size)
36
37
                     / Storage_Unit;
                  ~~~~~~~~~
38
          - -
          - -
39
                   Moving to next component
       end loop;
40
41
       for I in Arr'Range loop
42
         Put_Line ("Arr ("
43
                    & Integer'Image (I)
44
                    & ") :"
45
                    & Integer'Image (Arr (I)));
46
       end loop;
47
   end Show_Address;
48
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Addresses.Pointer_

⇔Arith_Ada

MD5: 2c1cdd6874036fb9a527baae63a312d9
```

Runtime output

Curr'Address : 00007FFF3744F690 Curr'Address : 00007FFF3744F694 Curr'Address : 00007FFF3744F698 Curr'Address : 00007FFF3744F690 Curr'Address : 00007FFF3744F6A0 Curr'Address : 00007FFF3744F6A4 Curr'Address : 00007FFF3744F6A8 Curr'Address : 00007FFF3744F6A0 Curr'Address : 00007FFF3744F6B0 Curr'Address : 00007FFF3744F6B0 Curr'Address : 00007FFF3744F6B4 Arr (1) : 1 Arr (2) : 2 Arr (3) : 3 Arr (4) : 4 Arr (5) : 5

(continues on next page)

Arr (6) : 6 Arr (7) : 7 Arr (8) : 8 Arr (9) : 9 Arr (10) : 10

1 2

3

4

5

6 7

8

9 10

11

12

13

14

15 16

17

18

19

20 21

22

23

In this example, we initialize the address A by retrieving the address of the first component of the array Arr. (Note that we could have written Arr(1) 'Address instead of Arr'Address. In any case, the language guarantees that Arr'Address gives us the address of the first component, i.e. Arr'Address = Arr(1) 'Address.)

Then, in the loop, we declare an overlay Curr using the current value of the A address. We can then operate on this overlay — here, we assign I to Curr. Finally, in the loop, we increment address A and make it *point* to the next component in the Arr array — to do so, we calculate the size of an **Integer** component in storage units. (For details on storage units, see the section on *storage size attribute* (page 85).)

```
In other languages
The code example above corresponds (more or less) to the following C code:
                                Listing 74: main.c
#include <stdio.h>
int main(int argc, const char * argv[])
{
    int i;
    int arr[10];
    int *a = arr;
    /* int *a = &arr[0]; */
    for (i = 0; i < 10; i++)
    {
        *a++ = i;
        printf("curr address: %p\n", a);
    }
    for (i = 0; i < 10; i++)
    {
        printf("arr[%d]: %d\n", i, arr[i]);
    }
    return 0;
}
Code block metadata
Project: Courses.Advanced Ada.Data Types.Type Representation.Addresses.Pointer
 →Arith C
MD5: 7aa709a4d7ed6ce2346dbabc853e28c0
Runtime output
```

```
curr address: 0x7ffddaedb0b4
curr address: 0x7ffddaedb0b8
curr address: 0x7ffddaedb0bc
curr address: 0x7ffddaedb0c0
curr address: 0x7ffddaedb0c4
curr address: 0x7ffddaedb0c8
curr address: 0x7ffddaedb0cc
curr address: 0x7ffddaedb0d0
curr address: 0x7ffddaedb0d4
curr address: 0x7ffddaedb0d8
arr[0]: 0
arr[1]: 1
arr[2]: 2
arr[3]: 3
arr[4]: 4
arr[5]: 5
arr[6]: 6
arr[7]: 7
arr[8]: 8
arr[9]: 9
```

While pointer arithmetic is very common in C, using address arithmetic in Ada is far from common, and it should be only used when it's really necessary to do so.

In the Ada Reference Manual

- 13.3 Operational and Representation Attributes⁷⁰
- 13.7.1 The Package System.Storage_Elements⁷¹

2.12 Discarding names

As we know, we can use the Image attribute of a type to get a string associated with this type. This is useful for example when we want to display a user message for an enumeration type:

```
Listing 75: show_enumeration_image.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
   procedure Show_Enumeration_Image is
3
4
      type Months is
5
         (January, February, March, April,
6
         May, June, July, August, September,
7
         October, November, December);
8
9
      M : constant Months := January;
10
   begin
11
      Put_Line ("Month: "
12
                 & Months'Image (M));
13
   end Show_Enumeration_Image;
14
```

Code block metadata

⁷⁰ http://www.ada-auth.org/standards/22rm/html/RM-13-3.html

⁷¹ http://www.ada-auth.org/standards/22rm/html/RM-13-7-1.html

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Discarding_Names. →Enumeration_Image MD5: 3863c5e06641d96b59edb9e76daa7560

Runtime output

```
Month: JANUARY
```

This is similar to having this code:

Listing 76: show enumeration image.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Enumeration Image is
3
4
      type Months is
5
         (January, February, March, April,
6
         May, June, July, August, September,
7
         October, November, December);
8
a
      M : constant Months := January;
10
11
      function Months Image (M : Months)
12
                               return String is
13
      begin
14
         case M is
15
                           => return "JANUARY";
16
             when January
             when February => return "FEBRUARY";
17
                            => return "MARCH";
            when March
18
                            => return "APRIL";
            when April
19
            when May
                            => return "MAY":
20
                            => return "JUNE";
            when June
21
                            => return "JULY";
            when July
22
                            => return "AUGUST";
            when August
23
             when September => return "SEPTEMBER";
24
             when October => return "OCTOBER";
25
             when November => return "NOVEMBER";
26
             when December => return "DECEMBER";
27
         end case;
28
      end Months_Image;
29
30
   beain
31
      Put_Line ("Month: "
32
                 & Months Image (M);
33
   end Show Enumeration Image;
34
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Discarding_Names.

→Enumeration_Image

MD5: 2db86044d2045bd9d4c3998cca36d51c
```

Runtime output

Month: JANUARY

Here, the Months_Image function associates a string with each month of the Months enumeration. As expected, the compiler needs to store the strings used in the Months_Image function when compiling this code. Similarly, the compiler needs to store strings for the Months enumeration for the Image attribute.

Sometimes, we don't need to call the Image attribute for a type. In this case, we could

save some storage by eliminating the strings associated with the type. Here, we can use the Discard_Names aspect to request the compiler to reduce — as much as possible — the amount of storage used for storing names for this type. Let's see an example:

```
Listing 77: show_discard_names.adb
```

```
procedure Show Discard Names is
1
      pragma Warnings (Off, "is not referenced");
2
3
      type Months is
4
         (January, February, March, April,
5
         May, June, July, August, September,
6
         October, November, December)
7
        with Discard_Names;
8
9
      M : constant Months := January;
10
   begin
11
      null:
12
   end Show_Discard_Names;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Discarding_Names. →Discard_Names MD5: 7891caac459a4be2096d443ca3190036

In this example, the compiler attempts to not store strings associated with the Months type duration compilation.

Note that the Discard_Names aspect is available for enumerations, exceptions, and tagged types.

1 In the GNAT toolchain

If we add this statement to the Show_Discard_Names procedure above:

we see that the application displays "0" instead of "JANUARY". This is because GNAT doesn't store the strings associated with the Months type when we use the Discard_Names aspect for the Months type. (Therefore, the Months'Image attribute doesn't have that information.) Instead, the compiler uses the integer value of the enumeration, so that Months'Image returns the corresponding string for this integer value.

In the Ada Reference Manual

Aspect Discard_Names⁷²

⁷² http://www.ada-auth.org/standards/22rm/html/RM-C-5.html

SHARED VARIABLE CONTROL

Ada has built-in support for handling both volatile and atomic data. Let's start by discussing volatile objects.

0	In	the	Ada	Reference	e Manual
---	----	-----	-----	-----------	----------

C.6 Shared Variable Control⁷³

3.1 Volatile

A volatile⁷⁴ object can be described as an object in memory whose value may change between two consecutive memory accesses of a process A — even if process A itself hasn't changed the value. This situation may arise when an object in memory is being shared by multiple threads. For example, a thread *B* may modify the value of that object between two read accesses of a thread *A*. Another typical example is the one of memory-mapped I/O⁷⁵, where the hardware might be constantly changing the value of an object in memory.

Because the value of a volatile object may be constantly changing, a compiler cannot generate code to store the value of that object in a register and then use the value from the register in subsequent operations. Storing into a register is avoided because, if the value is stored there, it would be outdated if another process had changed the volatile object in the meantime. Instead, the compiler generates code in such a way that the process must read the value of the volatile object from memory for each access.

Let's look at a simple example:

Listing 1: show_volatile_object.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Volatile Object is
3
      Val : Long Float with Volatile;
4
   begin
5
      Val := 0.0;
6
      for I in 0 .. 999 loop
7
         Val := Val + 2.0 * Long Float (I);
8
      end loop;
9
10
      Put Line ("Val: " & Long Float'Image (Val));
11
   end Show Volatile Object;
12
```

Code block metadata

- 73 http://www.ada-auth.org/standards/22rm/html/RM-C-6.html
- ⁷⁴ https://en.wikipedia.org/wiki/Volatile_(computer_programming)
- ⁷⁵ https://en.wikipedia.org/wiki/Memory-mapped_I/O

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Volatile.Object_ →Ada MD5: aale276e64e69813bfc3e3ef39f3dd47

Runtime output

Val: 9.9900000000000E+05

In this example, Val has the Volatile aspect, which makes the object volatile. We can also use the Volatile aspect in type declarations. For example:

Listing 2: shared_var_types.ads

```
1 package Shared_Var_Types is
2
3 type Volatile_Long_Float is new
4 Long_Float with Volatile;
5
6 end Shared Var Types;
```

Listing 3: show volatile type.adb

```
with Ada.Text IO;
                            use Ada.Text I0;
1
   with Shared_Var_Types; use Shared_Var_Types;
2
3
   procedure Show_Volatile_Type is
4
      Val : Volatile_Long_Float;
5
   begin
6
      Val := 0.0;
7
      for I in 0 .. 999 loop
8
         Val := Val + 2.0 * Volatile_Long_Float (I);
9
      end loop;
10
11
      Put Line ("Val: "
12
                 & Volatile Long Float'Image (Val));
13
   end Show_Volatile_Type;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Volatile.Type
MD5: 0d31156d47b2edcfb94debd016c8bb87
```

Runtime output

```
Val: 9.990000000000E+05
```

Here, we're declaring a new type Volatile_Long_Float in the Shared_Var_Types package. This type is based on the Long_Float type and uses the Volatile aspect. Any object of this type is automatically volatile.

In addition to that, we can declare components of an array to be volatile. In this case, we can use the Volatile_Components aspect in the array declaration. For example:

Listing 4: show_volatile_array_components.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Show_Volatile_Array_Components is
Arr : array (1 .. 2) of Long_Float
with Volatile_Components;
begin
```

```
Arr := (others => 0.0);
7
8
      for I in 0 .. 999 loop
9
         Arr (1) := Arr (1) + 2.0 * Long_Float (I);
10
         Arr (2) := Arr (2) + 10.0 * Long_Float (I);
11
      end loop;
12
13
      Put_Line ("Arr (1): "
14
                 & Long_Float'Image (Arr (1)));
15
      Put_Line ("Arr (2): "
16
                 & Long Float'Image (Arr (2)));
17
   end Show_Volatile_Array_Components;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Volatile.Array_

Gomponents

MD5: 05b3ee20f08c5a85f5872727a61c148d
```

Runtime output

Arr (1): 9.9900000000000E+05 Arr (2): 4.995000000000E+06

Note that it's possible to use the Volatile aspect for the array declaration as well:

Listing 5: shared_var_types.ads

```
1 package Shared_Var_Types is
2
3 private
4 Arr : array (1 .. 2) of Long_Float
5 with Volatile;
6
7 end Shared_Var_Types;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Volatile.Array
MD5: c9b7b9f94f1fac295753c7e7b9426fb2
```

Note that, if the Volatile aspect is specified for an object, then the Volatile_Components aspect is also specified automatically — if it makes sense in the context, of course. In the example above, even though Volatile_Components isn't specified in the declaration of the Arr array, it's automatically set as well.

3.2 Independent

When you write code to access a single object in memory, you might actually be accessing multiple objects at once. For example, when you declare types that make use of representation clauses — as we've seen in previous sections —, you might be accessing multiple objects that are grouped together in a single storage unit. For example, if you have components A and B stored in the same storage unit, you cannot update A without actually writing (the same value) to B. Those objects aren't independently addressable because, in order to access one of them, we have to actually address multiple objects at once.

When an object is independently addressable, we call it an independent object. In this case, we make sure that, when accessing that object, we won't be simultaneously accessing

another object. As a consequence, this feature limits the way objects can be represented in memory, as we'll see next.

To indicate that an object is independent, we use the Independent aspect:

Listing 6: shared_var_types.ads

```
package Shared_Var_Types is
    I : Integer with Independent;
    end Shared_Var_Types;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Independent.Object
MD5: d90fef37584ca8802b8a3e3858c0095b

Similarly, we can use this aspect when declaring types:

Listing 7: shared_var_types.ads

```
package Shared_Var_Types is
1
2
      type Independent Boolean is new Boolean
3
        with Independent;
4
5
      type Flags is record
6
         F1 : Independent Boolean;
7
         F2 : Independent Boolean;
8
      end record;
9
10
  end Shared_Var_Types;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Independent.Type
MD5: 7bcbee5b73067149b14c4b1b061f803c

In this example, we're declaring the Independent_Boolean type and using it in the declaration of the Flag record type. Let's now derive the Flags type and use a representation clause for the derived type:

Listing 8: sl	hared var	_types-represer	ntation.ads

```
package Shared_Var_Types.Representation is
1
2
      type Rep_Flags is new Flags;
3
4
      for Rep Flags use record
5
         F1 at 0 range 0 .. 0;
6
         F2 at 0 range 1 .. 1;
7
         - -
                        ^ ERROR: start position of
8
         - -
                                   F2 is wrong!
9
                           ERROR: F1 and F2 share the
                ^
10
         - -
                                   same storage unit!
11
          - -
      end record;
12
13
   end Shared_Var_Types.Representation;
14
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Independent.Type
MD5: bb9d5badf33401660e7e20a7cd612dab

Build output

```
shared_var_types-representation.ads:6:26: error: size for independent "F1" must be_

_multiple of Storage_Unit
shared_var_types-representation.ads:7:21: error: position for independent "F2"_

_must be multiple of Storage_Unit
shared_var_types-representation.ads:7:26: error: size for independent "F2" must be_

_multiple of Storage_Unit

gprbuild: *** compilation phase failed
```

As you can see when trying to compile this example, the representation clause that we used for Rep_Flags isn't following these limitations:

- 1. The size of each independent component must be a multiple of a storage unit.
- 2. The start position of each independent component must be a multiple of a storage unit.

For example, for architectures that have a storage unit of one byte — such as standard desktop computers —, this means that the size and the position of independent components must be a multiple of a byte. Let's correct the issues in the code above by:

- setting the size of each independent component to correspond to Storage_Unit using a range between 0 and Storage_Unit - 1 —, and
- setting the start position to zero.

This is the corrected version:

```
Listing 9: shared_var_types-representation.ads
```

```
with System;
1
2
   package Shared Var Types.Representation is
3
4
      type Rep Flags is new Flags;
5
6
      for Rep_Flags use record
7
         F1 at 0 range 0 .. System.Storage_Unit - 1;
8
         F2 at 1 range 0 .. System.Storage_Unit - 1;
9
      end record;
10
11
   end Shared_Var_Types.Representation;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Independent.Type
MD5: ed57e57cd746698909a4f7ce40a29dfc

Note that the representation that we're now using for Rep_Flags is most likely the representation that the compiler would have chosen for this data type. We could, however, have added an empty storage unit between F1 and F2 — by simply writing F2 **at** 2 ...:

Listing 10: shared_var_types-representation.ads

```
with System;
package Shared_Var_Types.Representation is
type Rep_Flags is new Flags;
```

```
6
7 for Rep_Flags use record
8 F1 at 0 range 0 .. System.Storage_Unit - 1;
9 F2 at 2 range 0 .. System.Storage_Unit - 1;
10 end record;
11
12 end Shared_Var_Types.Representation;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Independent.Type
MD5: 71fedf8aac7c19bca1ba3b487efa9b17

As long as we follow the rules for independent objects, we're still allowed to use representation clauses that don't correspond to the one that the compiler might select.

For arrays, we can use the Independent_Components aspect:

Listing 11: shared_var_types.ads

```
package Shared_Var_Types is
Flags : array (1 .. 8) of Boolean
with Independent_Components;
end Shared_Var_Types;
```

Code block metadata

We've just seen in a previous example that some representation clauses might not work with objects and types that have the Independent aspect. The same restrictions apply when we use the Independent_Components aspect. For example, this aspect prevents that array components are packed when the Pack aspect is used. Let's discuss the following erroneous code example:

Listing 12: shared_var_types.ads

```
1 package Shared_Var_Types is
2
3 type Flags is
4 array (Positive range <>) of Boolean
5 with Independent_Components, Pack;
6
7 F : Flags (1 .. 8) with Size => 8;
8
9 end Shared_Var_Types;
```

Code block metadata

Build output

shared_var_types.ads:5:37: warning: cannot pack independent components (RM 13.2(7))
shared_var_types.ads:7:36: error: size for "F" too small, minimum allowed is 64
gprbuild: *** compilation phase failed

As expected, this code doesn't compile. Here, we can have either independent components, or packed components. We cannot have both at the same time because packed components aren't independently addressable. The compiler warns us that the Pack aspect won't have any effect on independent components. When we use the Size aspect in the declaration of F, we confirm this limitation. If we remove the Size aspect, however, the code is compiled successfully because the compiler ignores the Pack aspect and allocates a larger size for F:

Listing 13: shared var types.ads

```
1 package Shared_Var_Types is
2
3 type Flags is
4 array (Positive range <>) of Boolean
5 with Independent_Components, Pack;
6
7 end Shared_Var_Types;
```

Listing 14: show flags size.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with System;
2
3
   with Shared Var Types; use Shared Var Types;
4
5
   procedure Show Flags Size is
6
      F : Flags (1 .. 8);
7
   begin
8
      Put Line ("Flags'Size:
9
                 & F'Size'Image & " bits");
10
      Put_Line ("Flags (1)'Size: "
11
                 & F (1) 'Size'Image & " bits");
12
      Put_Line ("# storage units:
13
                 & Integer'Image
14
                      (F'Size /
15
                      System.Storage Unit));
16
   end Show Flags Size;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Independent.

→Packed_Independent_Components

MD5: b96f921b08b1d8207749517f833fc121
```

Build output

Runtime output

Flags'Size: 64 bits Flags (1)'Size: 8 bits # storage units: 8

As you can see in the output of the application, even though we specify the Pack aspect for the Flags type, the compiler allocates eight storage units, one per each component of the

F array.

3.3 Atomic

An atomic object is an object that only accepts atomic reads and updates. The Ada standard specifies that "for an atomic object (including an atomic component), all reads and updates of the object as a whole are indivisible." In this case, the compiler must generate Assembly code in such a way that reads and updates of an atomic object must be done in a single instruction, so that no other instruction could execute on that same object before the read or update completes.

In other contexts

Generally, we can say that operations are said to be atomic when they can be completed without interruptions. This is an important requirement when we're performing operations on objects in memory that are shared between multiple processes.

This definition of atomicity above is used, for example, when implementing databases. However, for this section, we're using the term "atomic" differently. Here, it really means that reads and updates must be performed with a single Assembly instruction.

For example, if we have a 32-bit object composed of four 8-bit bytes, the compiler cannot generate code to read or update the object using four 8-bit store / load instructions, or even two 16-bit store / load instructions. In this case, in order to maintain atomicity, the compiler must generate code using one 32-bit store / load instruction.

Because of this strict definition, we might have objects for which the Atomic aspect cannot be specified. Lots of machines support integer types that are larger than the native word-sized integer. For example, a 16-bit machine probably supports both 16-bit and 32-bit integers, but only 16-bit integer objects can be marked as atomic — or, more generally, only objects that fit into at most 16 bits.

Atomicity may be important, for example, when dealing with shared hardware registers. In fact, for certain architectures, the hardware may require that memory-mapped registers are handled atomically. In Ada, we can use the Atomic aspect to indicate that an object is atomic. This is how we can use the aspect to declare a shared hardware register:

Listing 15: shared_var_types.ads

```
with System;
1
2
   package Shared_Var_Types is
3
4
   private
5
      R : Integer
6
             with Atomic,
7
                   Address =>
8
                     System'To Address (16#FFFF00A0#);
9
10
   end Shared_Var_Types;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic.Object
MD5: 5c2d8e0a9615084c2a15f896c61adaa6
```

Note that the **Address** aspect allows for assigning a variable to a specific location in the memory. In this example, we're using this aspect to specify the address of the memory-mapped register.

Later on, we talk again about the *Address aspect* (page 124) and the GNAT-specific *System*'*To_Address attribute* (page 125).

In addition to atomic objects, we can declare atomic types — similar to what we've seen before for volatile types. For example:

Listing 16: shared_var_types.ads

```
with System;
1
   package Shared Var Types is
3
4
      type Atomic_Integer is new Integer
5
        with Atomic:
6
7
   private
8
      R : Atomic_Integer
9
             with Address =>
10
                     System'To Address (16#FFFF00A0#);
11
12
   end Shared_Var_Types;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic.Types
MD5: 009632ba0155d70def8281ba590f3d12

In this example, we're declaring the Atomic_Integer type, which is an atomic type. Objects of this type — such as R in this example — are automatically atomic.

We can also declare atomic array components:

Listing 17: shared var types.ads

```
1 package Shared_Var_Types is
2
3 private
4 Arr : array (1 .. 2) of Integer
5 with Atomic_Components;
6
7 end Shared_Var_Types;
```

Code block metadata

This example shows the declaration of the Arr array, which has atomic components — the atomicity of its components is indicated by the Atomic_Components aspect.

Note that if an object is atomic, it is also volatile and independent. In other words, these type declarations are equivalent:

Listing 18: shared var types.ads

```
package Shared_Var_Types is
type Atomic_Integer_1 is new Integer
with Atomic;
type Atomic_Integer_2 is new Integer
with Atomic,
```

Volatile, Independent;

11 end Shared_Var_Types;

8

9 10

Code block metadata

A simular rule applies to components of an array. When we use the Atomic_Components, the following aspects are implied: Volatile, Volatile_Components and Independent_Components. For example, these array declarations are equivalent:

Listing 19: shared_var_types.ads

```
package Shared_Var_Types is
1
2
      Arr 1 : array (1 .. 2) of Integer
3
                 with Atomic Components;
4
5
      Arr_2 : array (1 .. 2) of Integer
6
                 with Atomic Components,
7
                      Volatile,
8
                      Volatile_Components,
9
                      Independent Components;
10
11
   end Shared Var Types;
12
```

3.4 Full-access only

1 Note

This feature was introduced in Ada 2022.

A full-access object is an object that requires that read or write operations on this object are performed by reading or writing all bits of the object (i.e. the *full object*) at once. Accordingly, a full-access type is a type whose objects follow this requirement. Note that a full-access type must be simultaneously a *volatile type* (page 135) or an *atomic type* (page 142). (In other words, if a type is neither volatile nor atomic, it cannot be a full-access type.)

1 Important

Just as a reminder, any atomic type is automatically also *volatile* (page 135) and *independent* (page 137).

Let's see some examples:

Listing 20: show_full_access_only_types.ads

```
package Show_Full_Access_Only_Types is
```

```
type Nonatomic_Full_Access_Type is
```

(continues on next page)

1

з

```
new Long_Float
4
           with Volatile, Full_Access_Only;
5
6
      type Atomic_Full_Access_Type is
7
         new Long_Float
8
           with Atomic, Full_Access_Only;
9
10
11
```

end Show_Full_Access_Only_Types;

Code block metadata

Project: Courses.Advanced Ada.Data Types.Type Representation.Shared Variable Gontrol.Full Access Only Types MD5: 6e7d4ee2e89b943d25319de9d8cebdcd

Likewise, we can define nonatomic and atomic full-access objects:

Listing 21: show_full_access_only_objects.ads

```
package Show_Full_Access_Only_Objects is
1
2
      Nonatomic Full Access Obj : Long Float
3
        with Volatile, Full_Access_Only;
4
5
      Atomic_Full_Access_Obj : Long_Float
6
        with Atomic, Full Access Only;
7
   end Show_Full_Access_Only_Objects;
```

Relevant topics

```
    9.10 Shared Variables<sup>76</sup>
```

C.6 Shared Variable Control⁷⁷

3.4.1 Nonatomic full-access

As we already know, the value of a volatile object may be constantly changing, so the compiler generates code to read the value of the volatile object from memory for each access. (In other words, the value cannot be stored in a register for further processing.)

In the case of nonatomic full-access objects, the value of the object must not only be read from memory or updated to memory every time, but those operations must also be performed for the complete record object — not just parts of it.

Consider the following example:

Listing 22: registers.ads

```
with System;
1
2
   package Registers is
3
4
      type Boolean Bit is new Boolean
5
        with Size => 1;
6
7
      type UInt1 is mod 2**1
8
```

(continues on next page)

⁷⁶ http://www.ada-auth.org/standards/22rm/html/RM-9-10.html ⁷⁷ http://www.ada-auth.org/standards/22rm/html/RM-C-6.html

```
with Size => 1;
9
10
      type UInt2 is mod 2**2
11
         with Size => 2;
12
13
      type UInt14 is mod 2**14
14
         with Size => 14;
15
16
      type Window_Register is record
17
          -- horizontal line count
18
          Horizontal_Cnt : UInt14 := 16#0#;
19
20
21
          -- unspecified
          Reserved_14_15 : UInt2 := 16#0#;
22
23
          -- vertical line count
24
          Vertical_Cnt : UInt14 := 16#0#;
25
26
          -- refresh signalling
27
          Refresh_Needed : Boolean_Bit := False;
28
29
          -- unspecified
30
          Reserved_30 : UInt1 := 16#0#;
31
      end record
32
33
         with Size
                        => 32,
              Bit_Order => System.Low_Order_First,
34
35
              Volatile,
              Full_Access_Only;
36
37
      for Window_Register use record
38
          Horizontal_Cnt at 0 range 0 .. 13;
39
          Reserved_14_15 at 0 range 14 .. 15;
40
          Vertical_Cnt at 0 range 16 .. 29;
41
          Refresh_Needed at 0 range 30 .. 30;
42
          Reserved_30
                        at 0 range 31 .. 31;
43
      end record;
44
45
      procedure Show (WR : Window_Register);
46
47
   end Registers;
48
```

```
Listing 23: registers.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Registers is
3
4
      procedure Show (WR : Window Register) is
5
      begin
6
          Put_Line ("WR = (Horizontal_Cnt => "
7
                     & WR.Horizontal Cnt'Image
8
                    & ",");
9
          Put_Line ("
                            Vertical Cnt
                                           => "
10
                    & WR.Vertical_Cnt'Image
11
          & ",");
Put_Line (" Refresh_Needed => "
12
13
                    & WR.Refresh Needed'Image
14
                    & ")");
15
      end Show;
16
17
   end Registers;
18
```

Code block metadata

In this example, we have a 32-bit register (of Window_Register type) that contains window information for a display:

position		0		
bits	#013	#14 #15	#16 #29	#30 #31
component	Horizontal_Cnt	Reserved_14_15	Vertical_Cnt	Reserved_30_31

Let's use the Window_Register type from the Registers package in a test application:

Listing 24: show_register.adb

```
with Registers;
                      use Registers;
1
2
   procedure Show Register is
3
      WR : Window_Register;
4
   begin
5
      -- Nonatomic full-access assignments
6
      WR.Horizontal_Cnt := 800;
7
      WR.Vertical_Cnt := 600;
8
      WR.Refresh Needed := True;
9
10
      Show (WR);
11
   end Show_Register;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Shared_Variable_

Gontrol.Nonatomic_Full_Access_Register

MD5: 7ff302d6cb282a6276747e8e17f26dfd
```

Runtime output

WR =	(Horizontal_Cnt	=>	800,
	Vertical_Cnt	=>	600,
	Refresh_Needed	=>	TRUE)

The example contains assignments such as WR.Horizontal_Cnt := 800 and WR. Vertical_Cnt:= 600. Because Window_Register is a full-access type, these assignments are performed for the complete 32-bit register, even though we're updating just a single component of the record object.

Note that if Window_Register wasn't a *full-access* object, an assignment such as WR. Horizontal_Cnt := 800 could be performed with a 16-bit operation. In fact, this is what a compiler would most probably select for this assignment, because that is more efficient than manipulating the entire object. Therefore, using a *full-access* object prevents the compiler from generating operations that could lead to unexpected results.

Whenever possible, this *full-access* assignment is performed in a single machine operation. However, if it's not possible to generate a single machine operation on the target machine, the compiler may generate multiple operations for the update of the record components.

Note that we could combine these two assignments into a single one using an aggregate:

Listing 25: show_register.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Registers;
                      use Registers;
3
4
   procedure Show Register is
5
      WR : Window Register;
6
   begin
7
      -- Nonatomic full-access assignment
8
      -- using an aggregate:
9
      WR := (Horizontal Cnt => 800,
10
              Vertical Cnt => 600,
11
              Refresh Needed => True,
12
              others
                             => <>);
13
14
      Show (WR);
15
  end Show Register;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Shared_Variable_

⊶Control.Nonatomic_Full_Access_Register

MD5: 9caf39e4a01ee1ec62f0b24747640c01
```

Runtime output

```
WR = (Horizontal_Cnt => 800,
    Vertical_Cnt => 600,
    Refresh_Needed => TRUE)
```

Again, this assignment is performed for the complete 32-bit register — ideally, using a single 32-bit machine operation — by reading the value from the memory.

Let's add another statement to the code example:

Listing 26: show_register.adb

```
with Registers; use Registers;
```

```
procedure Show_Register is
3
      WR : Window_Register :=
4
              (Horizontal_Cnt => 800,
5
               Vertical_Cnt => 600,
6
               Refresh_Needed => True,
7
               others
                               => <>):
8
   beain
9
      WR := (Horizontal_Cnt =>
10
                WR.Horizontal_Cnt * 2,
11
              Vertical_Cnt
                             =>
12
                Wr.Vertical_Cnt
                                   * 2,
13
              others
                              => <>);
14
15
      Show (WR);
16
   end Show_Register;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Shared_Variable_

Gontrol.Nonatomic_Full_Access_Register

MD5: cc4e218aef11af34e6d3262084a5c9ce
```

Runtime output

```
WR = (Horizontal_Cnt => 1600,
      Vertical_Cnt => 1200,
      Refresh_Needed => FALSE)
```

In this example, we have an initialization using the same aggregate as in the previous code example. We also have an assignment, in which we read the value of WR and use it in the calculation.

Delta aggregates

If we want to just change two components, but leave the information of other components untouched, we can use a *delta aggregate* (page 287). (Note that we haven't discussed the topic of delta aggregates yet: we'll do that *later on in this course* (page 287). However, in simple terms, we can use them to modify specific components of a record without changing the remaining components of the record.)

For example, we might want to update just the vertical count and indicate that update via the Refresh_Needed flag, but keep the same horizontal count:

Listing 27: show registers.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Registers;
                      use Registers;
3
4
   procedure Show_Registers is
5
      WR : Window Register :=
6
              (Horizontal_Cnt => 800,
7
               Vertical_Cnt
                             => 600,
8
               others
                               => <>);
9
   begin
10
          Delta assignment
       - -
11
      WR := (WR with delta
12
                   Vertical Cnt
                                   => 800,
13
                   Refresh Needed => True);
14
15
```

16 Show (WR); 17 end Show_Registers;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Shared_Variable_

Gontrol.Nonatomic_Full_Access_Register

MD5: 29df44d4fb13539cbd6070c37c217f8a
```

Runtime output

```
WR = (Horizontal_Cnt => 800,
      Vertical_Cnt => 800,
      Refresh_Needed => TRUE)
```

A delta assignment using an aggregate such as (WR with delta ...) includes reading the value of the complete 32-bit WR object from memory, changing the components specified after with delta, and writing the complete 32-bit WR object back to memory. The reason is that we need to retrieve the information that is supposed to remain intact — the Hori-zontal_Cnt and the reserved components — in order to write them back as a *full-access* operation.

3.4.2 Atomic full-access

As we already know, *atomic objects* (page 142) only accept atomic reads and updates, which — as a whole — are indivisible, i.e. they must be done in a single instruction, so that no other instruction could execute on that same object before the read or update completes. (Again, if an object is atomic, this implies it is also volatile.)

In the case of atomic full-access objects, the complete object must be read and updated. Ideally, this operation corresponds to a single atomic operation on the target machine, but it can also translate to multiple atomic operations.

Let's adapt the previous example to illustrate this. First, we adapt the type in the package:

Listing 28	registers.ads
------------	---------------

```
with System;
1
2
   package Registers is
3
4
      type Boolean Bit is new Boolean
5
        with Size => 1;
6
7
      type UInt1 is mod 2**1
8
        with Size => 1;
9
10
      type UInt2 is mod 2**2
11
        with Size => 2;
12
13
       type UInt14 is mod 2**14
14
        with Size => 14;
15
16
       type Window Register is record
17
          -- horizontal line count
18
          Horizontal Cnt : UInt14 := 16#0#;
19
20
          -- unspecified
21
          Reserved 14 15 : UInt2 := 16#0#;
22
23
```

```
-- vertical line count
24
          Vertical_Cnt : UInt14 := 16#0#;
25
26
          -- refresh signalling
27
          Refresh_Needed : Boolean_Bit := False;
28
29
          -- unspecified
30
          Reserved_30
                        : UInt1 := 16#0#;
31
       end record
32
        with Size
                        => 32,
33
              Bit_Order => System.Low_Order_First,
34
              Atomic,
35
              Full_Access_Only;
36
37
       for Window_Register use record
38
          Horizontal_Cnt at 0 range 0 .. 13;
39
          Reserved_14_15 at 0 range 14 .. 15;
40
          Vertical_Cnt at 0 range 16 .. 29;
41
          Refresh_Needed at 0 range 30 .. 30;
42
          Reserved_30
                         at 0 range 31 .. 31;
43
       end record;
44
45
      procedure Show (WR : Window_Register);
46
47
   end Registers;
48
```

Listing 29: registers.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Registers is
3
4
      procedure Show (WR : Window_Register) is
5
      begin
6
          Put_Line ("WR = (Horizontal Cnt => "
7
                    & WR.Horizontal_Cnt'Image
8
                    & ",");
9
                                           => "
          Put_Line ("
                            Vertical_Cnt
10
                    & WR.Vertical Cnt'Image
11
                    & ",");
12
          Put_Line ("
                           Refresh Needed => "
13
                    & WR.Refresh_Needed'Image
14
                    & ")");
15
      end Show;
16
17
   end Registers;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Shared_Variable_

GControl.Atomic_Full_Access_Register

MD5: dc088d1b0df1af5086a1ae8b46bb6d4d
```

We then use the package in our test application:

Listing 30: show_register.adb

```
with Registers; use Registers;
procedure Show_Register is
WR : Window_Register :=
```

```
(Horizontal_Cnt => 800,
5
               Vertical Cnt => 600,
6
               Refresh_Needed => True,
7
               others
                              => <>);
8
   beain
9
      WR := (Horizontal_Cnt =>
10
               WR.Horizontal_Cnt * 2,
11
              Vertical_Cnt =>
12
               Wr.Vertical_Cnt * 2,
13
                            => <>);
              others
14
15
      Show (WR);
16
17
   end Show_Register;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Type_Representation.Shared_Variable_

⊲Control.Atomic_Full_Access_Register

MD5: cc4e218aef11af34e6d3262084a5c9ce
```

In this example, we first have an atomic initialization of WR using an aggregate. Then, we have an atomic assignment to the atomic full-access object WR. Because its type is an atomic full-access type, the operations are atomic operations that always access the full object from and to memory.

3.4.3 Comparison: full-access and non-full-access types

An interesting exercise for the reader is to compare the Assembly code generated for the code example above with a version of this code where the Window_Register is not a full-access type.

Relevant topics

On a Linux platform, you can use *objdump* to retrieve the Assembly code and *diff* to see the difference between both versions of the type. For example:

```
objdump --target=elf64-x86-64 -d -S ./show_register > full_access.txt
```

```
# [...]
```

```
diff --width=80 -t -y full_access.txt no_full_access.txt
```

By doing this kind of comparisons, you might gain more insights on the impact of the Full_Access_Only aspect.

1 For further reading...

By running on a PC, we can compare the Intel Assembly⁷⁸ code for various versions of the code. Let's start with the version using a nonatomic full-access version of Window_Register vs. the nonatomic (non-full-access) version of Window_Register: type Window_Register is record -- [...] end record

```
with Size => 32,
Bit_Order => System.Low_Order_First,
Volatile,
Full_Access_Only;
```

type Window_Register is record [...] - end record with Size => 32, Bit_Order => System.Low_Order_First, Volatile; These are the manually-adapted differences between both versions: -- Volatile, Full Access Only | -- Volatile procedure Show_Register is procedure Show_Register is push %rbp push %rbp mov %rsp,%rbp mov %rsp,%rbp \$0x20,%rsp \$0x10,%rsp sub sub WR : Window_Register := WR : Window_Register := (Horizontal_Cnt => 800, (Horizontal_Cnt => 800, -0x4(%rbp),%eax -0x4(%rbp),%eax mov mov and \$0xffffc000,%eax and \$0xffffc000,%eax \$0x320,%eax or \$0x320,%eax or %eax,-0x4(%rbp) %eax,-0x4(%rbp) mov mov mov -0x4(%rbp),%eax mov -0x4(%rbp),%eax and \$0x3f,%ah and \$0x3f,%ah %eax,-0x4(%rbp) %eax,-0x4(%rbp) mov mov -0x4(%rbp),%eax -0x4(%rbp),%eax mov mov \$0xc000ffff,%eax and \$0xc000ffff,%eax and or \$0x2580000,%eax \$0x2580000,%eax or mov %eax,-0x4(%rbp) mov %eax,-0x4(%rbp) mov -0x4(%rbp),%eax mov -0x4(%rbp),%eax \$0x40000000,%eax \$0x4000000,%eax or or %eax,-0x4(%rbp) %eax,-0x4(%rbp) mov mov mov -0x4(%rbp),%eax mov -0x4(%rbp),%eax and \$0x7fffffff,%eax and \$0x7fffffff,%eax mov %eax,-0x4(%rbp) mov %eax,-0x4(%rbp) -0x4(%rbp),%eax mov < %eax,-0x14(%rbp) mov < -0x14(%rbp),%eax mov < mov %eax,-0x8(%rbp) < Vertical Cnt => 600, Vertical Cnt => 600, Refresh_Needed => True, Refresh_Needed => True, others => <>); others => <>): begin begin WR := (Horizontal_Cnt => WR := (Horizontal_Cnt => WR.Horizontal_Cnt * 2, WR.Horizontal Cnt * 2, -0x4(%rbp),%eax mov -0x8(%rbp),%eax mov %eax,%ecx mov < \$0x3fff,%cx and \$0x3fff,%ax and add %eax,%eax > WR := (Horizontal Cnt => < mov -0xc(%rbp),%eax < mov %eax,%edx < WR.Horizontal Cnt * 2, < lea (%rcx,%rcx,1),%eax < and \$0x3fff,%ax and \$0x3fff,%ax WR := (Horizontal_Cnt => WR := (Horizontal Cnt => movzwl %ax,%eax movzwl %ax,%eax \$0x3fff,%eax \$0x3fff,%eax and and \$0xffffc000,%edx and < %edx,%eax or < %eax,%edx mov %eax,%edx mov mov %edx,%eax mov -0x8(%rbp),%eax %eax,-0xc(%rbp) \$0xffffc000,%eax mov and

mov	-0xc(%rbp),%eax		 >	or mov	%edx,%eax %eax,-0x8(%rbp)	
			>	mov	-0x8(%rbp),%eax	
and	\$0x3f,%ah		-	and	\$0x3f,%ah	
mov	%eax,-0xc(%rbp)		1	mov	%eax,-0x8(%rbp)	
ino v	Vertical Cnt =>		I	110 0	Vertical Cnt =>	
	Wr.Vertical Cnt	* 2,			Wr.Vertical Cnt	* 2,
mov	-0x8(%rbp),%eax	Ζ,	1	mov	-0x4(%rbp),%eax	Ζ,
shr	\$0x10,%eax		I	shr	\$0x10,%eax	
	%eax,%ecx			and	\$0x3fff,%ax	
mov and	\$0x3fff,%cx			add	%eax,%eax	
				auu	sedx, sedx	
	(Horizontal_Cnt =>		< <			
mov	-Oxc(%rbp),%eax					
mov	%eax,%edx	* 2,	< <			
1	Wr.Vertical_Cnt					
lea	(%rcx,%rcx,1),%eax	(<			
and	\$0x3fff,%ax			and	\$0x3fff,%ax	
	(Horizontal_Cnt =>				(Horizontal_Cnt =>	
	l %ax,%eax				l %ax,%eax	
and	\$0x3fff,%eax			and	\$0x3fff,%eax	
shl	\$0x10,%eax			shl	\$0x10,%eax	
and	\$0xc000ffff,%edx		<			
or	%edx,%eax		<			
mov	%eax,%edx			mov	%eax,%edx	
mov	%edx,%eax			mov	-0x8(%rbp),%eax	
mov	%eax,-0xc(%rbp)			and	\$0xc000ffff,%eax	
mov	-0xc(%rbp),%eax			or	%edx,%eax	
			>	mov	%eax,-0x8(%rbp)	
			>	mov	-0x8(%rbp),%eax	
and	\$0xbfffffff,%eax			and	\$0xbfffffff,%eax	
mov	%eax,-0xc(%rbp)			mov	%eax,-0x8(%rbp)	
mov	-0xc(%rbp),%eax			mov	-0x8(%rbp),%eax	
and	\$0x7fffffff,%eax			and	\$0x7fffffff,%eax	
mov	%eax,-0xc(%rbp)		<			
mov	-0xc(%rbp),%eax		<			
mov	%eax,-0x8(%rbp)			mov	%eax,-0x8(%rbp)	
			>	mov	-0x8(%rbp),%eax	
			>	mov	%eax,-0x4(%rbp)	
	others => <	<>);			others => <	<>);

As we can see, although parts of the Assembly code are the same or look very similar, there are some differences between both versions. These differences are mostly related to the fact that we have to operate on the full object when reading it from memory.

Likewise, we can compare the Assembly code for the atomic full-access version of Window_Register vs. the atomic (non-full-access) version of Window_Register:

```
type Window_Register is record
  -- [...]
end record
with Size => 32,
  Bit_Order => System.Low_Order_First,
  Atomic,
  Full_Access_Only;
type Window_Register is record
  -- [...]
end record
with Size => 32,
  Bit_Order => System.Low_Order_First,
  Atomic;
```

These are the manually-adapted differences between both versions:

Atomic,	Full_Access_Only	I	Atomic	
procedure S push mov sub WR : Win (mov and or mov mov and or mov mov and or mov wr WR : Win mov mov wr WR : Win mov mov and wr WR : (H mov mov and WR := (H	<pre>how_Register is %rbp %rsp,%rbp \$0x20,%rsp dow_Register := Horizontal_Cnt => 800, -0x4(%rbp),%eax \$0xffffc000,%eax \$0x320,%eax %eax,-0x4(%rbp) -0x4(%rbp),%eax \$0x3f,%ah %eax,-0x4(%rbp) -0x4(%rbp),%eax \$0x2580000,%eax %eax,-0x4(%rbp) -0x4(%rbp),%eax \$0x2580000,%eax %eax,-0x4(%rbp) -0x4(%rbp),%eax %eax,-0x4(%rbp) -0x4(%rbp),%eax %eax,-0x4(%rbp) dow_Register := -0x4(%rbp),%eax %eax,-0x4(%rbp) -0x14(%rbp),%eax %eax,-0x8(%rbp) Vertical_Cnt => 600, Refresh_Needed => True others => <>); Orizontal_Cnt => WR.Horizontal_Cnt * 2 -0x8(%rbp),%eax %eax,%ecx \$0x3fff,%cx</pre>		<pre>procedure Show_Register is push %rbp mov %rsp,%rbp sub \$0x10,%rsp WR : Window_Register :=</pre>	, e , ;
<pre>mov mov mov mov wov wR := (H mov mov and WR := (H mov mov lea and WR := (H movzwl and</pre>	<pre>-0x4(%rbp),%eax %eax,-0x14(%rbp) -0x14(%rbp),%eax %eax,-0x8(%rbp) Vertical_Cnt => 600, Refresh_Needed => True others => <>); Vorizontal_Cnt => WR.Horizontal_Cnt * 2 -0x8(%rbp),%eax %eax,%ecx \$0x3fff,%cx Vorizontal_Cnt => -0xc(%rbp),%eax %eax,%edx WR.Horizontal_Cnt * 2 (%rcx,%rcx,1),%eax</pre>	<	<pre>mov -0x4(%rbp),%eax mov %eax,-0x8(%rbp) Vertical_Cnt => 600 Refresh_Needed => Tru others => <>) begin WR := (Horizontal_Cnt => WR.Horizontal_Cnt * mov -0x8(%rbp),%eax and \$0x3fff,%ax add %eax,%eax</pre>	e, ;
or mov mov mov and mov V	<pre>\$0xffffc000,%edx %edx,%eax %eax,%edx %edx,%eax %eax,-0xc(%rbp) -0xc(%rbp),%eax \$0x3f,%ah %eax,-0xc(%rbp) ertical_Cnt => Wr.Vertical_Cnt * 2 -0x8(%rbp),%eax</pre>	 < <	<pre>and \$0xffffc000,%eax or %edx,%eax mov %eax,-0xc(%rbp) mov -0xc(%rbp),%eax and \$0x3f,%ah mov %eax,-0xc(%rbp) Vertical_Cnt => Wr.Vertical_Cnt * mov -0x8(%rbp),%eax</pre>	2,

```
shr
           $0x10,%eax
                                               shr
                                                      $0x10,%eax
    mov
           %eax,%ecx
                                        <
           $0x3fff,%cx
                                        $0x3fff,%ax
    and
                                               and
                                               add
                                                      %eax,%eax
                                        >
   WR := (Horizontal_Cnt =>
                                        <
           -0xc(%rbp),%eax
    mov
                                        <
           %eax,%edx
    mov
                                        <
                               * 2,
            Wr.Vertical Cnt
                                        <
    lea
           (%rcx,%rcx,1),%eax
                                        <
           $0x3fff,%ax
                                               and
                                                      $0x3fff,%ax
    and
   WR := (Horizontal Cnt =>
                                              WR := (Horizontal Cnt =>
    movzwl %ax,%eax
                                               movzwl %ax,%eax
                                                      $0x3fff,%eax
    and
           $0x3fff,%eax
                                               and
    shl
           $0x10,%eax
                                               shl
                                                      $0x10,%eax
                                               mov
                                        >
                                                      %eax,%edx
                                        >
                                               mov
                                                      -Oxc(%rbp),%eax
    and
           $0xc000ffff,%edx
                                        and
                                                      $0xc000ffff,%eax
    or
           %edx.%eax
                                               or
                                                      %edx,%eax
           %eax,%edx
    mov
                                        <
           %edx,%eax
    mov
                                        <
           %eax,-0xc(%rbp)
                                                      %eax,-0xc(%rbp)
    mov
                                               mov
    mov
           -Oxc(%rbp),%eax
                                               mov
                                                      -Oxc(%rbp),%eax
    and
           $0xbfffffff,%eax
                                               and
                                                      $0xbfffffff,%eax
    mov
           %eax,-0xc(%rbp)
                                               mov
                                                      %eax,-0xc(%rbp)
           -Oxc(%rbp),%eax
                                                      -Oxc(%rbp),%eax
    mov
                                               mov
                                               and
           $0x7fffffff,%eax
                                                      $0x7ffffff,%eax
    and
                                               mov
    mov
           %eax,-0xc(%rbp)
                                                      %eax,-0xc(%rbp)
           -0xc(%rbp),%eax
                                                      -Oxc(%rbp),%eax
    mov
                                               mov
    xcha
           %eax,-0x8(%rbp)
                                               xchg
                                                      %eax,-0x8(%rbp)
          others
                         => <>);
                                                     others
                                                                     => <>);
Again, there are some differences between both versions, even though some parts of
the Assembly code are the same or look very similar.
Finally, we might want to compare the nonatomic full-access version vs. the atomic
full-access version of the Window Register type:
type Window Register is record
       [...]
end record
  with Size
                 => 32,
       Bit Order => System.Low Order First,
       Volatile,
       Full Access Only;
type Window Register is record
   -- [...]
end record
  with Size
                 => 32,
       Bit_Order => System.Low_Order_First,
       Atomic.
       Full Access Only;
These are the differences between both versions:
-- Volatile, Full_Access_Only
                                      | -- Atomic, Full_Access_Only
procedure Show_Register is
                                           procedure Show_Register is
    push
           %rbp
                                               push
                                                      %rbp
    mov
           %rsp,%rbp
                                               mov
                                                      %rsp,%rbp
    sub
           $0x20,%rsp
                                               sub
                                                      $0x20,%rsp
   WR : Window_Register :=
                                              WR : Window_Register :=
          (Horizontal_Cnt => 800,
                                                     (Horizontal_Cnt => 800,
                                                      -0x4(%rbp),%eax
    mov
           -0x4(%rbp),%eax
                                               mov
```

\$0xffffc000,%eax

and

and

\$0xffffc000,%eax

\$0x320,%eax or mov %eax,-0x4(%rbp) mov -0x4(%rbp),%eax and \$0x3f,%ah %eax,-0x4(%rbp) mov mov -0x4(%rbp),%eax and \$0xc000ffff,%eax \$0x2580000.%eax or mov %eax,-0x4(%rbp) mov -0x4(%rbp),%eax\$0x4000000,%eax or %eax,-0x4(%rbp) mov mov -0x4(%rbp),%eax \$0x7ffffff,%eax and mov %eax,-0x4(%rbp) WR : Window_Register := -0x4(%rbp),%eax mov mov %eax,-0x14(%rbp) mov -0x14(%rbp),%eax mov %eax,-0x8(%rbp) => 600, Vertical Cnt Refresh Needed => True, others => <>); begin WR := (Horizontal Cnt => WR.Horizontal_Cnt * 2, mov -0x8(%rbp),%eax %eax,%ecx mov \$0x3fff,%cx and WR := (Horizontal_Cnt => mov -0xc(%rbp),%eax mov %eax,%edx WR.Horizontal Cnt * 2, lea (%rcx,%rcx,1),%eax and \$0x3fff,%ax WR := (Horizontal_Cnt => movzwl %ax,%eax \$0x3fff,%eax and \$0xffffc000,%edx and or %edx,%eax %eax,%edx mov mov %edx,%eax %eax,-0xc(%rbp) mov mov -0xc(%rbp),%eax and \$0x3f,%ah mov %eax,-0xc(%rbp) Vertical_Cnt => * 2, Wr.Vertical_Cnt mov -0x8(%rbp),%eax shr \$0x10,%eax mov %eax,%ecx and \$0x3fff,%cx WR := (Horizontal Cnt => mov -Oxc(%rbp),%eax mov %eax,%edx Wr.Vertical Cnt * 2. lea (%rcx,%rcx,1),%eax \$0x3fff,%ax and WR := (Horizontal_Cnt => movzwl %ax,%eax and \$0x3fff,%eax shl \$0x10,%eax

```
$0x320,%eax
    or
    mov
           %eax,-0x4(%rbp)
    mov
           -0x4(%rbp),%eax
    and
           $0x3f,%ah
    mov
           %eax,-0x4(%rbp)
           -0x4(%rbp),%eax
    mov
    and
           $0xc000ffff,%eax
           $0x2580000,%eax
    or
    mov
           %eax,-0x4(%rbp)
    mov
            -0x4(%rbp),%eax
           $0x4000000,%eax
    or
           %eax,-0x4(%rbp)
    mov
    mov
            -0x4(%rbp),%eax
           $0x7ffffff,%eax
    and
           %eax,-0x4(%rbp)
    mov
   WR : Window_Register :=
            -0x4(%rbp),%eax
    mov
    mov
           %eax,-0x14(%rbp)
    mov
            -0x14(%rbp),%eax
    mov
           %eax,-0x8(%rbp)
                           => 600.
           Vertical Cnt
           Refresh Needed => True,
           others
                           => <>);
begin
   WR := (Horizontal Cnt =>
            WR.Horizontal_Cnt * 2,
    mov
           -0x8(%rbp),%eax
           %eax,%ecx
    mov
           $0x3fff,%cx
    and
   WR := (Horizontal_Cnt =>
    mov
            -0xc(%rbp),%eax
    mov
           %eax,%edx
            WR.Horizontal Cnt * 2,
    lea
            (%rcx,%rcx,1),%eax
    and
           $0x3fff,%ax
   WR := (Horizontal_Cnt =>
    movzwl %ax,%eax
           $0x3fff,%eax
    and
           $0xffffc000,%edx
    and
    or
           %edx,%eax
    mov
           %eax,%edx
    mov
           %edx,%eax
           %eax,-0xc(%rbp)
    mov
    mov
            -0xc(%rbp),%eax
    and
           $0x3f,%ah
    mov
           %eax,-0xc(%rbp)
          Vertical_Cnt
                          =>
                                * 2,
            Wr.Vertical_Cnt
    mov
            -0x8(%rbp),%eax
    shr
           $0x10,%eax
    mov
           %eax,%ecx
    and
           $0x3fff,%cx
   WR := (Horizontal_Cnt =>
    mov
            -0xc(%rbp),%eax
    mov
           %eax,%edx
            Wr.Vertical Cnt
                                * 2,
    lea
            (%rcx,%rcx,1),%eax
    and
           $0x3fff,%ax
   WR := (Horizontal_Cnt =>
    movzwl %ax,%eax
    and
           $0x3fff,%eax
    shl
           $0x10,%eax
```

and	\$0xc000ffff,%edx		and	\$0xc000ffff,%edx
or	%edx,%eax		or	%edx,%eax
mov	%eax,%edx		mov	%eax,%edx
mov	%edx,%eax		mov	%edx,%eax
mov	%eax,-0xc(%rbp)		mov	%eax,-0xc(%rbp)
mov	-0xc(%rbp),%eax		mov	-0xc(%rbp),%eax
and	\$0xbfffffff,%eax		and	\$0xbfffffff,%eax
mov	%eax,-0xc(%rbp)		mov	%eax,-0xc(%rbp)
mov	-0xc(%rbp),%eax		mov	-0xc(%rbp),%eax
and	\$0x7fffffff,%eax		and	\$0x7fffffff,%eax
mov	%eax,-0xc(%rbp)		mov	%eax,-0xc(%rbp)
mov	-0xc(%rbp),%eax		mov	-0xc(%rbp),%eax
mov	%eax,-0x8(%rbp)		xchg	%eax,-0x8(%rbp)
	others => <>);			others => <>);

As we can see, the code is basically the same — except for the last Assembly instruction, which is a *mov* instruction in the volatile version and an *xchg* instruction in the atomic version — which is an atomic instruction on this platform.

3.5 Atomic operations

\rm Note

This feature was introduced in Ada 2022.

Ada offers four packages to handle atomic operations. Those packages are child packages of the System.Atomic_Operations package. We will discuss each of those package individually in this section.

Relevant topics

C.6.1 The Package System.Atomic_Operations⁷⁹

3.5.1 Atomic Exchange

The generic System.Atomic_Operations.Exchange package provides operations to compare and exchange objects atomically.

Atomic_Exchange function

One of those operations is the Atomic_Exchange function, which performs the following operations atomically:

```
function Atomic_Exchange
 (Item : aliased in out Atomic_Type;
  Value : Atomic_Type)
  return Atomic_Type
is
   Old_Item : Atomic_Type := Item;
begin
   Item := Value;
```

⁷⁸ https://en.wikipedia.org/wiki/X86_instruction_listings

⁷⁹ http://www.ada-auth.org/standards/22rm/html/RM-C-6-1.html

return Old_Item; end Atomic_Exchange;

As mentioned in the Ada Reference Manual⁸⁰, we can use this function to implement a spinlock⁸¹. For example:

Listing 31: spinlocks.ads

```
with System.Atomic Operations.Exchange;
1
2
   package Spinlocks is
3
4
      type Lock is new Boolean with Atomic;
5
6
      package Lock Exchange is new
7
         System.Atomic_Operations.Exchange (Lock);
8
9
   end Spinlocks;
10
```

Listing 32: show_locks.adb

```
with Spinlocks;
1
   use Spinlocks;
2
   use Spinlocks.Lock_Exchange;
3
4
   procedure Show_Locks is
5
      L : aliased Lock := False;
6
   beain
7
          Get the lock
8
      while Atomic_Exchange (Item => L,
9
                               Value => True) loop
10
         null;
11
12
      end loop;
13
       -- At this point, we got the lock.
14
       -- Do some stuff here...
15
16
       -- Release the lock.
17
      L := False:
18
   end Show Locks;
19
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations. →Exchange MD5: 36699b917485f14c4e8a905a6c48027b

In this example, we call the Atomic_Exchange function for the L lock until we get it. Then, we can use the resource that we protected via the lock. After we finish our work, we can release the lock by setting L to **False**.

Note that System.Atomic_Operations.Exchange is a generic package, so we have to instantiate it for a specific atomic type — in this case, the atomic Boolean Lock type.

We can use multiple tasks to illustrate a situation where using a lock is important to ensure that no race conditions⁸² occur:

⁸⁰ http://www.ada-auth.org/standards/22rm/html/RM-C-6-2.html

⁸¹ https://en.wikipedia.org/wiki/Spinlock

⁸² https://en.wikipedia.org/wiki/Race_condition

Listing 33: spinlocks.ads

```
with System.Atomic_Operations.Exchange;
1
2
   package Spinlocks is
3
4
      type Lock is new Boolean with Atomic;
5
6
      package Lock_Exchange is new
7
        System.Atomic_Operations.Exchange (Lock);
8
9
   end Spinlocks;
10
```

Listing 34: show locks.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Spinlocks;
3
   use Spinlocks;
4
   use Spinlocks.Lock_Exchange;
5
6
   procedure Show_Locks is
7
                  : aliased Lock := False;
       L
8
       Task_Count : Integer := 0;
9
10
11
       task type A_Task;
12
       task body A_Task is
13
          Task_Number : Integer;
14
       begin
15
             Get the lock
          - -
16
          while Atomic_Exchange (Item => L,
17
                                   Value => True) loop
18
             null;
19
          end loop;
20
21
          -- At this point, we got the lock.
22
          Task_Count := Task_Count + 1;
23
          Task Number := Task Count;
24
25
          -- Release the lock.
26
          L := False;
27
28
          Put_Line ("Task_Number: "
29
                     & Task Number'Image);
30
31
32
       end A_Task;
33
       A, B, C, D, E, F : A_Task;
34
   begin
35
       null;
36
   end Show_Locks;
37
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations. →Exchange MD5: af7aad741c20be1e8433b04def90dcdb

Runtime output

Task_Number: 1 Task_Number: 6 Task_Number: 2 Task_Number: 5 Task_Number: 3 Task_Number: 4

In this example, we create multiple tasks (A, B, C, D, E, F) and initialize the **Task_Number** of each task based on the value of the **Task_Count** variable. To avoid multiple tasks accessing the **Task_Count** variable at the same time, we use the L lock, which we get before updating the **Task_Count**.

Atomic_Compare_And_Exchange function

Another function from the System.Atomic_Operations.Exchange package is Atomic_Compare_And_Exchange, which performs the following operations atomically:

```
function Atomic Compare And Exchange
        aliased in out Atomic_Type;
  (Item
  Prior
          : aliased in out Atomic Type;
  Desired :
                            Atomic Type
  return Boolean is
begin
   if Item = Prior then
      Item := Value:
       -- The item is only updated if its
       -- value and the prior value match
       return True;
   else
      Prior := Item;
      return False;
   end if;
end Atomic Exchange;
```

This function can be used for lazy initialization⁸³ of variables. For example, consider an application with multiple tasks that make use of a certain value that isn't initialized at its declaration, but at a later point in time by an arbitrary task. We can use Atomic_Compare_And_Exchange to ensure that we only update that value if it wasn't already initialized.

Let's start with the package specification:

```
Listing 35: lazy_initialization.ads
```

```
with System.Atomic Operations.Exchange;
1
   with Ada.Numerics.Discrete_Random;
2
3
   package Lazy Initialization is
4
5
      subtype Lazy Value Total Range is
6
        Integer range 99 .. 1000;
7
8
      Lazy Value Default Value : constant
9
          := Lazy_Value_Total_Range'First;
10
11
      subtype Lazy_Value_Range is Integer
12
         range Lazy Value Default_Value + 1 ..
13
               Lazy_Value_Total_Range'Last;
14
```

(continues on next page)

⁸³ https://en.wikipedia.org/wiki/Lazy_initialization

```
15
       type Lazy_Value is new Lazy_Value_Total_Range
16
         with Atomic,
17
              Default_Value =>
18
                Lazy_Value_Default_Value;
19
20
       package Value_Exchange is new
21
         System.Atomic_Operations.Exchange
22
           (Lazy_Value);
23
24
      package Lazy Value Random is new
25
         Ada.Numerics.Discrete_Random
26
           (Lazy_Value_Range);
27
28
   end Lazy_Initialization;
29
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations. →Compare_And_Exchange MD5: 09d49998aa7e3d5c0cfb4b74af8e542b

In this package, we declare the Lazy_Value type with a default value (specified by the Lazy_Value_Default_Value constant). Note that we have two ranges here: Lazy_Value_Total_Range and Lazy_Value_Range. We use the Lazy_Value_Total_Range in the declaration of the Lazy_Value type: it indicates the *total range* of the type. We use the Lazy_Value_Range as a constraint for the total range. This range doesn't contain the default value (Lazy_Value_Default_Value), and we use it to indicate the valid values of the type. (We discuss the application of Lazy_Value_Range later on.)

Also, in addition to instantiating the System.Atomic_Operations.Exchange package, we instantiate the Ada.Numerics.Discrete_Random package, which we'll use to generate random numbers in the expected range (Lazy_Value_Range) for the Lazy_Value type. (We discussed the Ada.Numerics.Discrete_Random package in the Introduction to Ada⁸⁴ course.)

Let's use this package in the Show_Lazy_Initialization procedure:

Listing 36	: show	lazy	initialization.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   with Ada.Numerics.Discrete Random;
2
3
   with Lazy_Initialization;
4
   use Lazy Initialization;
5
   use Lazy_Initialization.Value_Exchange;
6
   procedure Show_Lazy_Initialization is
8
     subtype A Task Number is Natural;
9
10
      Value
                          : aliased Lazy Value;
11
      Value Modified By : A Task Number := 0;
12
13
      task type A_Task is
14
          entry Start (This : A_Task_Number);
15
          entry Stop;
16
      end A_Task;
17
18
      task body A Task is
19
```

(continues on next page)

⁸⁴ https://learn.adacore.com/courses/intro-to-ada/chapters/standard_library_numerics.html# intro-ada-random-number-generation

```
Task_Number : A_Task_Number;
20
      begin
21
          accept Start (This : A_Task_Number) do
22
             Task_Number := This;
23
          end Start;
24
25
          Sleep_Some_Time : declare
26
             subtype Sleep_Range is
27
               Integer range 1 .. 3;
28
29
             package Random_Sleep is new
30
               Ada.Numerics.Discrete_Random
31
                  (Sleep_Range);
32
             use Random_Sleep;
33
             G : Generator;
35
          begin
36
             Reset (G);
37
             delay Duration (Random (G));
38
          end Sleep_Some_Time;
39
40
          Generate_Value : declare
41
             use Lazy_Value_Random;
42
43
             G
                                        Generator;
44
             Initial_Value :
45
                                        Lazy_Value_Range;
46
             Prior
                             : aliased Lazy_Value;
          begin
47
             Reset (G);
48
             Initial_Value := Random (G);
49
50
             if Atomic_Compare_And_Exchange
51
                         => Value,
                (Item
52
                       => Prior,
                Prior
53
                Desired => Lazy_Value (Initial_Value))
54
             then
55
                Value_Modified_By := Task_Number;
56
             end if;
57
58
          end Generate_Value;
59
60
          accept Stop do
61
             Put_Line ("Current task number:
62
                        & Task_Number'Image);
63
             Put Line ("Value:
64
                        & Value'Image);
65
             Put_Line ("Modified by task number: "
66
                        & Value_Modified_By'Image);
67
             Put_Line ("-----");
68
          end Stop;
69
      end A_Task;
70
71
      Some_Tasks : array (1 .. 5) of A_Task;
72
   begin
73
       for I in Some Tasks'Range loop
74
          Some_Tasks (I).Start (I);
75
       end loop;
76
       for I in Some Tasks'Range loop
77
          Some_Tasks (I).Stop;
78
       end loop;
79
   end Show_Lazy_Initialization;
80
```

34

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations. →Compare_And_Exchange MD5: 9cc898edd767f8bbcfe2c81e7ca0e442

Runtime output

Current task number: 1 Value: 539 Modified by task number: 1 Current task number: 2 Value: 539 Modified by task number: 1 -----Current task number: 3 Value: 539 Modified by task number: 1 Current task number: 4 Value: 539 Modified by task number: 1 Current task number: 5 Value: 539 Modified by task number: 1

In the Show_Lazy_Initialization procedure, the most important variable is Value, which is the variable we have to protect via a lock. In addition, we have the auxiliary Value_Modified_By variable, which indicates the number of the task that initialized the Value variable.

In this procedure, we also see two main *block statements* (page 460):

- the block statement with the Sleep_Some_Time identifier, where we make the task *sleep* for a random amount of time (in the Sleep_Range range); and
- the block statement with the Generate_Value identified, where we generate a new value randomly and attempt to update the Value variable (of Lazy_Value type).

Let's discuss some details about the Generate_Value block statement. We start by declaring some variables. Here, it's important to highlight that the Prior variable is initialized with the default value (Lazy_Value_Default_Value). We then call the Atomic_Compare_And_Exchange function, and pass Value and Prior as actual parameters. We can have two possible outcomes:

- If Value hasn't been modified by a task yet, it will contain the default value which means that the values of the Prior and Value variables match. In this case, the call to Atomic_Compare_And_Exchange will update the Value variable and return True. (Note that we also update the Value_Modified_By variable when Atomic_Compare_And_Exchange returns True.)
- 2. If Value has already been modified by a task, its value doesn't match the (default) value of Prior anymore, so the call to Atomic_Compare_And_Exchange doesn't modify the Value variable.

As mentioned before, we use a stricter range for the random number generator: the Lazy_Value_Range. Because this range doesn't contain the default value (Lazy_Value_Default_Value), we will never generate a random value that matches the default value.

Relevant topics

C.6.2 The Package System.Atomic_Operations.Exchange⁸⁵

Atomic Test and Set

The System.Atomic_Operations.Test_And_Set package provides atomic operations to set and clear atomic flags. To declare flags, we use the Test_And_Set_Flag type. The following operations are available:

- 1. the Atomic_Test_And_Set function, which we call to verify whether the flag can be set and, if positive, set it accordingly.
 - The function returns **True** if the flag has been set, and **False** otherwise.
- 2. the Atomic_Clear procedure, which we call to clear the flag.

We can use these functions to implement an application similar to the *spinlocks* (page 159) that we've seen before:

Listing 37: show_test_and_set.adb

```
with System. Atomic Operations. Test And Set;
1
   use System.Atomic_Operations.Test_And_Set;
2
3
   with Ada.Text IO; use Ada.Text IO;
4
5
   procedure Show_Test_And_Set is
6
                 : aliased Test_And_Set_Flag;
       Lock
7
       Task_Count : Integer := 0;
8
9
       task type A_Task;
10
11
       task body A_Task is
12
          Task_Number : Integer;
13
       begin
14
              Get the lock
15
          while Atomic_Test_And_Set (Lock) loop
16
17
             null;
          end loop;
18
19
          -- At this point, we got the lock.
20
          Task Count := Task Count + 1;
21
          Task_Number := Task_Count;
22
23
          -- Release the lock.
24
          Atomic Clear (Lock);
25
26
          Put_Line ("Task_Number: "
27
                     & Task Number'Image);
28
29
       end A_Task;
30
31
       A, B, C, D, E, F : A_Task;
32
   begin
33
       null;
34
   end Show Test And Set;
35
```

Code block metadata

⁸⁵ http://www.ada-auth.org/standards/22rm/html/RM-C-6-2.html

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations. ⇔Test_And_Set MD5: 45814e2e157d3fd45f876c89914a5cc5

Runtime output

Task_Number: 1 Task_Number: 4 Task_Number: 5 Task_Number: 3 Task_Number: 6 Task_Number: 2

Here, we call Atomic_Test_And_Set in a loop until it returns **True**. Then, we update the **Task_Count** and **Task_**Number. When we're finished, we call the Atomic_Clear procedure to release the lock.

Relevant topics

C.6.3 The Package System.Atomic_Operations.Test_and_Set⁸⁶

Atomic Operations using Integer Arithmetic

The generic System.Atomic_Operations.Integer_Arithmetic package is used to perform atomic operations on atomic integer types. It provides the following operations: the procedures Atomic_Add and Atomic_Subtract, and the functions Atomic_Fetch_And_Add and Atomic_Fetch_And_Subtract. The procedures and the corresponding Atomic_Fetch_ functions do basically the same thing, with the difference that Atomic_Fetch functions return the previous (older) value of the input item.

The Atomic Add procedure performs the following operations atomically:

```
procedure Atomic_Add
 (Item : aliased in out Atomic_Type;
   Value : Atomic_Type) is
   begin
        Item := Item + Value;
end Atomic_Add;
```

The corresponding Atomic_Fetch_And_Add function performs the following operations atomically:

```
function Atomic_Fetch_And_Add
 (Item : aliased in out Atomic_Type;
   Value : Atomic_Type
   is
        Old_Item : Atomic_Type := Item;
begin
        Item := Item + Value;
        return Old_Item;
end Atomic_Fetch_And_Add;
```

The Atomic Subtract procedure performs the following operations atomically:

procedure Atomic_Subtract
 (Item : aliased in out Atomic_Type;

⁸⁶ http://www.ada-auth.org/standards/22rm/html/RM-C-6-3.html

```
Value : Atomic_Type) is
begin
Item := Item - Value;
end Atomic_Subtract;
```

The corresponding Atomic_Fetch_And_Subtract function performs the following operations atomically:

```
function Atomic_Fetch_And_Subtract
 (Item : aliased in out Atomic_Type;
 Value : Atomic_Type
    return Atomic_Type
    is
        Old_Item : Atomic_Type := Item;
    begin
        Item := Item - Value;
        return Old_Item;
end Atomic_Fetch_And_Subtract;
```

Let's reuse a *previous code example* (page 159) that sets a unique number for each task. In this case, instead of using locks, we use the atomic operations from the System. Atomic_Operations.Integer_Arithmetic package:

Listing 38: atomic_integers.ads

```
with System.Atomic_Operations.Integer_Arithmetic;
1
2
   package Atomic_Integers is
3
4
      type Atomic Integer is new Integer
5
        with Atomic;
6
7
      package Atomic Integer Arithmetic is new
8
         System.Atomic_Operations.Integer_Arithmetic
9
           (Atomic_Integer);
10
11
   end Atomic_Integers;
12
```

Listing 39: show_atomic_integers.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Atomic_Integers;
3
   use Atomic_Integers;
4
   use Atomic_Integers.Atomic_Integer_Arithmetic;
5
6
   procedure Show_Atomic_Integers is
7
      Task_Count : aliased Atomic_Integer := 0;
8
9
      task type A_Task;
10
11
      task body A Task is
12
          Task_Number : Atomic_Integer;
13
      begin
14
15
          Task Number :=
16
            Atomic_Fetch_And_Add (Task_Count, 1);
17
          Put_Line ("Task_Number: "
18
                    & Task_Number'Image);
19
20
```

```
21 end A_Task;
22 A, B, C, D, E, F : A_Task;
24 begin
25 null;
26 end Show_Atomic_Integers;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations.

→Integer_Arithmetic

MD5: 835093f90b9efe37b93ca84fe1ce3444
```

Runtime output

Task_Number: 0 Task_Number: 1 Task_Number: 4 Task_Number: 3 Task_Number: 5 Task_Number: 2

In this example, we call the Atomic_Fetch_And_Add function to update the Task_Count variable and, at the same time, initialize the Task_Number variable of the current task.

Relevant topics

• C.6.4 The Package System. Atomic_Operations. Integer_Arithmetic⁸⁷

Atomic Operations using Modular Arithmetic

The generic System.Atomic_Operations.Modular_Arithmetic package is very similar to the System.Atomic_Operations.Integer_Arithmetic package. In fact, it provides the same operations: the procedures Atomic_Add and Atomic_Subtract, and the functions Atomic_Fetch_And_Add and Atomic_Fetch_And_Subtract. The only difference is that it is used for modular types instead of integer types.

Let's reuse the *previous code example* (page 167), but replace the atomic integer type by an atomic modular type:

	Listing 40	atomic	modu	lars.ads
--	------------	--------	------	----------

```
with System.Atomic_Operations.Modular_Arithmetic;
1
2
   package Atomic_Modulars is
3
4
      type Atomic Modular is mod 100
5
        with Atomic;
6
7
      package Atomic_Modular_Arithmetic is new
8
         System.Atomic_Operations.Modular_Arithmetic
9
           (Atomic_Modular);
10
11
   end Atomic Modulars;
12
```

⁸⁷ http://www.ada-auth.org/standards/22rm/html/RM-C-6-4.html

Listing 41: show_atomic_modulars.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Atomic Modulars;
3
   use Atomic Modulars;
4
   use Atomic Modulars.Atomic Modular Arithmetic;
5
6
   procedure Show_Atomic_Modulars is
7
      Task_Count : aliased Atomic_Modular := 0;
8
9
      task type A_Task;
10
11
      task body A_Task is
12
          Task Number : Atomic Modular;
13
      begin
14
          Task Number :=
15
            Atomic Fetch And Add (Task Count, 1);
16
17
          Put_Line ("Task_Number: "
18
                    & Task_Number'Image);
19
20
      end A_Task;
21
22
      A, B, C, D, E, F : A_Task;
23
   begin
24
      null;
25
   end Show Atomic Modulars;
26
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Shared_Variable_Control.Atomic_Operations. ⊶Modular_Arithmetic MD5: 3a5a85febacd13f5e053cf00b19746ff

Runtime output

Task_Number: 0 Task_Number: 1 Task_Number: 4 Task_Number: 2 Task_Number: 5 Task_Number: 3

As we did in the previous example, we again call the Atomic_Fetch_And_Add function to update the Task_Count variable and, at the same time, initialize the Task_Number variable of the current task. The only difference is that we use a modular type (Atomic_Modular).

Relevant topics

• C.6.5 The Package System. Atomic_Operations. Modular_Arithmetic⁸⁸

⁸⁸ http://www.ada-auth.org/standards/22rm/html/RM-C-6-5.html

CHAPTER FOUR

RECORDS

4.1 Default Initialization

As mentioned in the Introduction to Ada⁸⁹ course, record components can have default initial values. Also, we've seen that other kinds of types can have *default values* (page 65).

In the Ada Reference Manual, we refer to these default initial values as "default expressions of record components." The term *default expression* indicates that we can use any kind of expression for the default initialization of record components — which includes subprogram calls for example:

Listing 1: show_default_initialization.ads

```
package Show Default Initialization is
1
2
      function Init return Integer is
3
         (42);
4
5
      type Rec is record
6
         A : Integer := Init;
7
      end record;
8
9
  end Show_Default_Initialization;
10
```

Code block metadata

In this example, the A component is initialized by default by a call to the Init procedure.

1 In the Ada Reference Manual

```
• 3.8 Record Types<sup>90</sup>
```

4.1.1 Dependencies

Default expressions cannot depend on other components. For example, if we have two components A and B, we cannot initialize B based on the value that A has:

⁸⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/records.html#intro-ada-record-default-values ⁹⁰ http://www.ada-auth.org/standards/22rm/html/RM-3-8.html

Listing 2: show_default_initialization_dependency.ads

```
package Show Default Initialization Dependency is
1
2
      function Init return Integer is
3
        (42);
4
5
      type Rec is record
6
         A : Integer := Init;
7
         B : Integer := Rec.A; -- Illegal!
8
      end record;
9
10
   end Show Default Initialization Dependency;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Default_Initialization.No_

→Dependency

MD5: ca23cbd7e4a54d0b9c6974aed0ee77c8
```

Build output

In this example, we cannot initialize the B component based on the value of the A component. (In fact, the syntax Rec.A as a way to refer to the A component is only allowed in predicates, not in the record component declaration.)

4.1.2 Initialization Order

The default initialization of record components is performed in arbitrary order. In fact, the order is decided by the compiler, so we don't have control over it.

Let's see an example:

Listing 3: simple_recs.ads

```
package Simple_Recs is
1
2
      function Init (S : String;
3
                       I : Integer)
4
                       return Integer;
5
6
      type Rec is record
7
          A : Integer := Init ("A", 1);
8
          B : Integer := Init ("B", 2);
9
      end record;
10
11
   end Simple Recs;
12
```

Listing 4: simple_recs.adb

```
with Ada.Text_I0; use Ada.Text_I0;
package body Simple_Recs is
function Init (S : String;
I : Integer)
return Integer is
```

```
8 begin
9 Put_Line (S & ": " & I'Image);
10 return I;
11 end Init;
12
13 end Simple_Recs;
```

Listing 5: show_initialization_order.adb

```
with Simple_Recs; use Simple_Recs;
procedure Show_Initialization_Order is
R : Rec;
begin
null;
r end Show_Initialization_Order;
```

Code block metadata

Runtime output

A: 1 B: 2

When running this code example, you might see this:

A: 1 B: 2

However, the compiler is allowed to rearrange the operations, so this output is possible as well:

B: 2

A: 1

Therefore, we must write the default expression of each individual record components in such a way that the resulting initialization value is always correct, independently of the order that those expressions are evaluated.

4.1.3 Evaluation

According to the Annotated Ada Reference Manual, the "default expression of a record component is only evaluated upon the creation of a default-initialized object of the record type." This means that the default expression is by itself not evaluated when we declare the record type, but when we create an object of this type. It follows from this rule that the default is only evaluated when necessary, i.e,, when an explicit initial value is not specified in the object declaration.

Let's see an example:

Listing 6: show_initialization_order.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Simple_Recs; use Simple_Recs;
```

```
procedure Show_Initialization_Order is
4
   begin
5
      Put_Line ("Some processing first...");
6
      Put_Line
7
         ("Now, let's declare an object "
8
         & "of the record type Rec...");
9
10
      declare
11
         R : Rec;
12
      begin
13
          Put_Line
14
            ("An object of Rec type has "
15
             & "just been created.");
16
      end;
17
18
   end Show_Initialization_Order;
19
```

Code block metadata

Runtime output

```
Some processing first...
Now, let's declare an object of the record type Rec...
A: 1
B: 2
An object of Rec type has just been created.
```

Here, we only see the information displayed by the Init function — which is called to initialize the A and B components of the R record — during the object creation. In other words, the default expressions Init ("A", 1) and Init ("B", 2) are *not* evaluated when we declare the R type, but when we create an object of this type.

In the Ada Reference Manual

```
    3.8 Record Types<sup>91</sup>
```

4.1.4 Defaults and object declaration

Note

This subsection was originally written by Robert A. Duff and published as Gem #12: Limited Types in Ada 2005⁹².

Consider the following type declaration:

Listing 7: type_defaults.ads

```
package Type_Defaults is
    type Color_Enum is (Red, Blue, Green);
```

(continues on next page)

```
    <sup>91</sup> http://www.ada-auth.org/standards/22aarm/html/AA-3-8.html
    <sup>92</sup> https://www.adacore.com/gems/ada-gem-12
```

1

2 3

```
type T is private;
4
   private
5
      type T is
6
         record
7
                      : Color_Enum := Red;
             Color
8
             Is_Gnarly : Boolean := False;
9
             Count
                       : Natural;
10
         end record;
11
12
      procedure Do_Something;
13
  end Type_Defaults;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Default_Initialization.Default_

↓Init

MD5: 218154278081f89595534bc02e34539b
```

If we want to say, "make **Count** equal 100, but initialize Color and Is_Gnarly to their defaults", we can do this:

Listing 8: type_defaults.adb

```
package body Type_Defaults is
1
2
      Object 100 : constant T :=
3
                      (Color => <>,
4
                       Is Gnarly => <>,
5
                                 => 100);
                       Count
6
7
      procedure Do Something is null;
8
9
  end Type Defaults;
10
```

Code block metadata

Historically

1 2

3

4

5

6 7

8 9 Prior to Ada 2005, the following style was common:

```
Listing 9: type_defaults.adb
```

```
package body Type_Defaults is
```

```
Object_100 : constant T :=
    (Color => Red,
    Is_Gnarly => False,
    Count => 100);
procedure Do_Something is null;
```

10 end Type_Defaults;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Default_Initialization.Default_

→Init

MD5: clddfae75d7f0c691356027903a6d144
```

Here, we only wanted Object_100 to be a default-initialized T, with **Count** equal to 100. It's a little bit annoying that we had to write the default values Red and **False** twice. What if we change our mind about Red, and forget to change it in all the relevant places? Since Ada 2005, the <> notation comes to the rescue, as we've just seen.

On the other hand, if we want to say, "make **Count** equal 100, but initialize all other components, including the ones we might add next week, to their defaults", we can do this:

Listing 10: type_defaults.adb

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Default_Initialization.Default_

⊲Init

MD5: 93f5d71ae80ff0ebad54f2569539f536
```

Note that if we add a component Glorp : **Integer**; to type T, then the **others** case leaves Glorp undefined just as this code would do:

Listing 11: type_defaults.adb

```
package body Type_Defaults is
procedure Do_Something is
Object_100 : T;
begin
Object_100.Count := 100;
end Do_Something;
end Type_Defaults;
```

Code block metadata

Therefore, you should be careful and think twice before using others.

4.1.5 Advanced Usages

In addition to expressions such as subprogram calls, we can use *per-object expressions* (page 243) for the default value of a record component. (We discuss this topic later on in more details.)

For example:

Listing 12: rec_per_object_expressions.ads

```
package Rec_Per_Object_Expressions is
```

```
type T (D : Positive) is private;
3
4
   private
5
6
      type T (D : Positive) is record
7
          V : Natural := D - 1;
8
9
          - -
                Per-object expression
          - -
10
      end record;
11
12
   end Rec Per Object Expressions;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Default_Initialization.Per_Object_

→Expressions

MD5: 92591ea482db2b009b8eeafe633ca6cd
```

In this example, component V is initialized by default with the per-object expression D $\,$ - $\,$ 1, where D refers to the discriminant D.

4.2 Mutually dependent types

In this section, we discuss how to use *incomplete types* (page 36) to declare mutually dependent types. Let's start with this example:

Listing 13:	mutually	_dependent.ads

```
package Mutually Dependent is
1
2
       type T1 is record
3
          B : T2;
4
       end record;
5
6
       type T2 is record
7
         A : T1;
8
      end record;
9
10
   end Mutually Dependent;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Mutually_Dependent_Types.Mutually_

→Dependent

MD5: ffa8d6ab83a1172dcbae0978952dacb2
```

Build output

mutually_dependent.ads:4:11: error: "T2" is undefined
gprbuild: *** compilation phase failed

When you try to compile this example, you get a compilation error. The first problem with this code is that, in the declaration of the T1 record, the compiler doesn't know anything about T2. We could solve this by declaring an incomplete type (**type T2**;) before the declaration of T1. This, however, doesn't solve all the problems in the code: the compiler still doesn't know the size of T2, so we cannot create a component of this type. We could, instead, declare an access type and use it here. By doing this, even though the compiler doesn't know the size of T2, it knows the size of an access type designating T2, so the record component can be of such an access type.

To summarize, in order to solve the compilation error above, we need to:

- use at least one incomplete type;
- declare at least one component as an access to an object.

For example, we could declare an incomplete type T2 and then declare the component B of the T1 record as an access to T2. This is the corrected version:

```
Listing 14: mutually_dependent.ads
```

```
package Mutually_Dependent is
1
2
      type T2;
3
      type T2_Access is access T2;
4
5
       type T1 is record
6
         B : T2 Access;
7
      end record;
8
9
      type T2 is record
10
         A : T1;
11
      end record;
12
13
   end Mutually_Dependent;
14
```

Code block metadata

We could strive for consistency and declare two incomplete types and two accesses, but this isn't strictly necessary in this case. Here's the adapted code:

Listing 15: mutually dependent.ads

```
package Mutually_Dependent is
1
2
      type T1;
3
      type T1_Access is access T1;
4
5
      type T2;
6
      type T2_Access is access T2;
7
8
9
       type T1 is record
         B : T2_Access;
10
      end record;
11
12
       type T2 is record
13
         A : T1 Access;
14
       end record;
15
16
  end Mutually_Dependent;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Mutually_Dependent_Types.Mutually_ →Dependent MD5: 9a9899cd0dd2525bd27d67d6629a0071

Later on, we'll see that these code examples can be written using *anonymous access types* (page 734).

```
In the Ada Reference Manual
```

3.10.1 Incomplete Type Declarations⁹³

4.3 Null records

A null record is a record that doesn't have any components. Consequently, it cannot store any information. When declaring a null record, we simply write **null** instead of declaring actual components, as we usually do for records. For example:

Listing 16: null recs.ads

```
1 package Null_Recs is
2
3 type Null_Record is record
4 null;
5 end record;
6
7 end Null_Recs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Null_Record
MD5: 3c82da822710342354134fa71a03452a
```

Note that the syntax can be simplified to **is null record**, which is much more common than the previous form:

Listing 17: null recs.ads

```
package Null_Recs is
type Null_Record is null record;
end Null_Recs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Null_Record
MD5: 1da1746ce5b0a237276272d2b620e282
```

Although a null record doesn't have components, we can still specify subprograms for it. For example, we could specify an addition operation for it:

Listing 18: null_recs.ads

```
1 package Null_Recs is
2
3 type Null_Record is null record;
4
5 function "+" (A, B : Null_Record)
6 return Null_Record;
7
8 end Null_Recs;
```

⁹³ http://www.ada-auth.org/standards/22rm/html/RM-3-10-1.html

Listing 19: null_recs.adb

```
package body Null_Recs is
1
2
      function "+" (A, B : Null_Record)
3
                      return Null_Record
4
      is
5
         pragma Unreferenced (A, B);
6
      begin
7
         return (null record);
8
      end "+";
9
10
   end Null_Recs;
11
```

Listing 20: show_null_rec.adb

```
with Null_Recs; use Null_Recs;

procedure Show_Null_Rec is
    A, B : Null_Record;

begin
    B := A + A;
    A := A + B;
end Show Null_Rec;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Null_Record
MD5: 3a1c2fbae75541dfb0b2ff4c14d22039

1 In the Ada Reference Manual

4.3.1 Record Aggregates⁹⁴

4.3.1 Simple Prototyping

A null record doesn't provide much functionality on itself, as we're not storing any information in it. However, it's far from being useless. For example, we can make use of null records to design an API, which we can then use in an application without having to implement the actual functionality of the API. This allows us to design a prototype without having to think about all the implementation details of the API in the first stage.

Consider this example:

```
Listing 21: devices.ads
```

```
package Devices is
1
2
      type Device is private;
3
4
      function Create
5
        (Active : Boolean)
6
         return Device;
7
8
      procedure Reset
9
         (D : out Device) is null;
10
11
```

(continues on next page)

⁹⁴ http://www.ada-auth.org/standards/22rm/html/RM-4-3-1.html

```
procedure Process
12
         (D : in out Device) is null;
13
14
       procedure Activate
15
         (D : in out Device) is null;
16
17
       procedure Deactivate
18
         (D : in out Device) is null;
19
20
   private
21
22
       type Device is null record;
23
24
       function Create (Active : Boolean)
25
                          return Device is
26
         (null record);
27
28
   end Devices;
29
```

Listing 22: show_device.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Devices;
                     use Devices;
2
З
   procedure Show_Device is
4
      A : Device;
5
   begin
6
      Put Line ("Creating device...");
7
      A := Create (Active => True);
8
9
      Put Line ("Processing on device...");
10
      Process (A);
11
12
      Put_Line ("Deactivating device...");
13
      Deactivate (A);
14
15
      Put Line ("Activating device...");
16
      Activate (A);
17
18
      Put_Line ("Resetting device...");
19
      Reset (A);
20
   end Show_Device;
21
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Device
MD5: 7d2fce20ac33607f7081381b307a564a

Runtime output

```
Creating device...
Processing on device...
Deactivating device...
Activating device...
Resetting device...
```

In the Devices package, we're declaring the Device type and its primitive subprograms: Create, Reset, Process, Activate and Deactivate. This is the API that we use in our prototype. Note that, although the Device type is declared as a private type, it's still defined as a null record in the full view.

In this example, the Create function, implemented as an expression function in the private

part, simply returns a null record. As expected, this null record returned by Create matches the definition of the Device type.

All procedures associated with the Device type are implemented as null procedures, which means they don't actually have an implementation nor have any effect. We'll discuss this topic *later on in the course* (page 500).

In the Show_Device procedure — which is an application that implements our prototype —, we declare an object of Device type and call all subprograms associated with that type.

4.3.2 Extending the prototype

Because we're either using expression functions or null procedures in the specification of the Devices package, we don't have a package body for it (as there's nothing to be implemented). We could, however, move those user messages from the Show_Devices procedure to a dummy implementation of the Devices package. This is the adapted code:

```
package Devices is
1
2
      type Device is null record;
3
4
      function Create (Active : Boolean)
5
                         return Device;
6
7
      procedure Reset (D : out Device);
8
9
      procedure Process (D : in out Device);
10
11
      procedure Activate (D : in out Device);
12
13
      procedure Deactivate (D : in out Device);
14
15
   end Devices;
16
```

Listing 24: devices.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Devices is
3
4
       function Create (Active : Boolean)
5
                         return Device
6
      is
7
          pragma Unreferenced (Active);
8
      begin
9
          Put Line ("Creating device...");
10
          return (null record);
11
      end Create:
12
13
      procedure Reset (D : out Device)
14
       is
15
          pragma Unreferenced (D);
16
      begin
17
          Put Line ("Processing on device...");
18
      end Reset;
19
20
      procedure Process (D : in out Device)
21
       is
22
          pragma Unreferenced (D);
23
      begin
24
```

```
Put_Line ("Deactivating device...");
25
       end Process;
26
27
       procedure Activate (D : in out Device)
28
29
       is
          pragma Unreferenced (D);
30
      beain
31
          Put_Line ("Activating device...");
32
      end Activate;
33
34
      procedure Deactivate (D : in out Device)
35
       is
36
37
          pragma Unreferenced (D);
38
      begin
          Put_Line ("Resetting device...");
39
       end Deactivate;
40
41
   end Devices;
42
```

Listing 25: show_device.adb

```
with Devices; use Devices;
1
2
   procedure Show Device is
З
      A : Device;
4
   begin
5
      A := Create (Active => True);
6
      Process (A);
7
8
      Deactivate (A);
9
      Activate (A);
      Reset (A);
10
   end Show_Device;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Device
MD5: 1a21b41f3847f6c132ccbc9696ab7689

Runtime output

Creating device... Deactivating device... Resetting device... Activating device... Processing on device...

As we changed the specification of the Devices package to not use null procedures, we now need a corresponding package body for it. In this package body, we implement the operations on the Device type, which actually just display a user message indicating which operation is being called.

Let's focus on this updated version of the Show_Device procedure. Now that we've removed all those calls to Put_Line from this procedure and just have the calls to operations associated with the Device type, it becomes more apparent that, even though Device is just a null record, we can design an application with a sequence of various commands operating on it. Also, when we just read the source-code of the Show_Device procedure, there's no clear indication that the Device type doesn't actually hold any information.

4.3.3 More complex applications

As we've just seen, we can use null records like any other type and create complex prototypes with them. We could, for instance, design an application that makes use of many null records, or even have types that depend on or derive from null records. Let's see a simple example:

```
Listing 26: many_devices.ads
```

```
package Many_Devices is
1
2
      type Device is null record;
3
4
      type Device Config is null record;
5
6
      function Create (Config : Device_Config)
7
                         return Device is
8
         (null record);
9
10
      type Derived Device is new Device;
11
12
      procedure Process (D : Derived Device) is null;
13
14
   end Many_Devices;
15
```

Listing 27: show_derived_device.adb

```
with Many_Devices; use Many_Devices;
1
2
   procedure Show_Derived_Device is
3
      A : Device;
4
      B : Derived Device;
5
      C : Device_Config;
6
   begin
7
      A := Create (Config => C);
8
      B := Create (Config => C);
9
10
      Process (B);
11
   end Show Derived Device;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Derived_Device
MD5: 757a3def24c8333a27b64943727d8d4e

In this example, the Create function has a null record parameter (of Device_Config type) and returns a null record (of Device type). Also, we derive the Derived_Device type from the Device type. Consequently, Derived_Device is also a null record (since it's derived from a null record). In the Show_Derived_Device procedure, we declare objects of those types (A, B and C) and call primitive subprograms to operate on them.

This example shows that, even though the types we've declared are *just* null records, they can still be used to represent dependencies in our application.

4.3.4 Implementing the API

Let's focus again on the previous example. After we have an initial prototype, we can start implementing some of the functionality needed for the Device type. For example, we can store information about the current activation state in the record:

```
Listing 28: devices.ads
```

```
package Devices is
1
2
       type Device is private;
3
4
       function Create (Active : Boolean)
5
                         return Device;
6
7
       procedure Reset (D : out Device);
8
9
       procedure Process (D : in out Device);
10
11
       procedure Activate (D : in out Device);
12
13
       procedure Deactivate (D : in out Device);
14
15
   private
16
17
       type Device is record
18
          Active : Boolean;
19
       end record;
20
21
   end Devices;
22
```

```
Listing 29: devices.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Devices is
3
4
      function Create (Active : Boolean)
5
                         return Device
6
      is
7
          pragma Unreferenced (Active);
8
9
      begin
          Put_Line ("Creating device...");
10
          return (Active => Active);
11
      end Create;
12
13
      procedure Reset (D : out Device)
14
      is
15
          pragma Unreferenced (D);
16
      begin
17
          Put Line ("Processing on device...");
18
      end Reset;
19
20
      procedure Process (D : in out Device)
21
      is
22
          pragma Unreferenced (D);
23
      begin
24
          Put_Line ("Deactivating device...");
25
      end Process;
26
27
      procedure Activate (D : in out Device)
28
       is
29
      begin
30
          Put_Line ("Activating device...");
31
          D.Active := True;
32
      end Activate;
33
34
```

```
35 procedure Deactivate (D : in out Device)
36 is
37 begin
38 Put_Line ("Resetting device...");
39 D.Active := False;
40 end Deactivate;
41
42 end Devices;
```

Listing 30: show_device.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Devices;
                    use Devices:
2
3
   procedure Show Device is
4
      A : Device;
5
   begin
6
      A := Create (Active => True);
7
      Process (A);
8
      Deactivate (A);
9
      Activate (A);
10
      Reset (A);
11
  end Show_Device;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Device
MD5: 348ce0c110b47a6b6fd1c9fe73ef0558

Build output

```
devices.adb:11:25: warning: aspect Unreferenced specified for "Active" [enabled by

→default]
```

Runtime output

```
Creating device...
Deactivating device...
Resetting device...
Activating device...
Processing on device...
```

Now, the Device record contains an Active component, which is used in the updated versions of Create, Activate and Deactivate.

Note that we haven't done any change to the implementation of the Show_Device procedure: it's still the same application as before. As we've been hinting in the beginning, using null records makes it easy for us to first create a prototype — as we did in the Show_Device procedure — and postpone the API implementation to a later phase of the project.

4.3.5 Tagged null records

A null record may be tagged, as we can see in this example:

Listing 31: null_recs.ads

```
1 package Null_Recs is
2
3 type Tagged_Null_Record is
4 tagged null record;
```

```
5
6 type Abstract_Tagged_Null_Record is
7 abstract tagged null record;
8
9 end Null_Recs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Tagged_Null_Record
MD5: 918572d2c50911b84c80a9c601b75439
```

As we see in this example, a type can be **tagged**, or even **abstract tagged**. We discuss abstract types later on in the course.

As expected, in addition to deriving from tagged types, we can also extend them. For example:

Listing 32: devices.ads

```
package Devices is
1
2
3
      type Device is private;
4
       function Create (Active : Boolean)
5
                         return Device;
6
7
      type Derived Device is private;
8
9
   private
10
11
       type Device is tagged null record;
12
13
       function Create (Active : Boolean)
14
                         return Device is
15
         (null record);
16
17
       type Derived Device is new Device with record
18
          Active : Boolean;
19
       end record;
20
21
       function Create (Active : Boolean)
22
                         return Derived_Device is
23
         (Active => Active);
24
25
   end Devices:
26
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Null_Records.Extended_Device
MD5: 15e06a5115cbcb131477b5224a6594db

In this example, we derive Derived_Device from the Device type and extend it with the Active component. (Because we have a type extension, we also need to override the Create function.)

Since we're now introducing elements from object-oriented programming, we could consider using interfaces instead of null records. We'll discuss this topic later on in the course.

4.4 Record discriminants

We introduced the topic of record discriminants in the Introduction to Ada course⁹⁵. Also, in a previous chapter, we mentioned that record types with unconstrained discriminants without defaults are *indefinite types* (page 31).

In this section, we discuss a couple of details about record discriminants that we haven't covered yet. Although the discussion will be restricted to record discriminants, keep in mind that tasks and protected types can also have discriminants. We'll focus on discriminants for tasks and protected types in separate chapters.

In addition, discriminants can be used to write *per-object expressions* (page 240). We discuss this topic later in this chapter.

1 In the Ada Reference Manual

• 3.7 Discriminants⁹⁶

4.4.1 Known and unknown discriminant parts

When it comes to discriminants, a type declaration falls into one of the following three categories: it has either no discriminants at all, known discriminants or unknown discriminants.

In order to have no discriminants, a type simply doesn't have a discriminant part in its declaration. For example:

Listing 33:	show	discriminants.ads
-------------	------	-------------------

```
package Show Discriminants is
1
2
      type T No Discr is private;
3
4
       - -
          no discriminant part
5
6
   private
7
8
       type T No Discr is null record;
9
10
   end Show Discriminants;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.No_Discriminant_Part
MD5: f6701bd9c46b265753a258a6f99a5c7a

By using parentheses after the type name, we're defining a discriminant part. In this case, the type can either have unknown or known discriminants. For example:

Listing 34: show_discriminants.ads

```
package Show_Discriminants is

type T_Unknown_Discr (<>) is

...
Unknown discriminant
private;
```

⁹⁵ https://learn.adacore.com/courses/intro-to-ada/chapters/more_about_records.html# intro-ada-record-discriminants

⁹⁶ http://www.ada-auth.org/standards/12rm/html/RM-3-7.html

```
7
       type T_Known_Discr (D : Integer) is
8
9
            Known discriminant
10
         private;
11
12
   private
13
14
       type T_Unknown_Discr is
15
         null record;
16
17
       type T_Known_Discr (D : Integer) is
18
         null record;
19
20
   end Show_Discriminants;
21
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Discriminant_Parts
MD5: 486edc81b72473e022bb9e56ebaca559

An unknown discriminant part is represented by (<>) in the partial view — this is basically the so-called *box notation* <> (also known as *box compound delimiter*) in parentheses. We discuss unknown discriminant parts and their peculiarities *later on in this chapter* (page 217). In this section, we mainly focus on known discriminants.

We've already seen examples of known discriminants in previous chapters. In simple terms, known discriminants are composed by one or more discriminant specifications, which are similar to subprogram parameters, but without parameter modes. In fact, we can think of discriminants as parameters for a type T, but with the goal of defining specific characteristics or constraints when declaring objects of type T.

4.4.2 Discriminant as constant property

We can think of discriminants as constant properties of a type. In fact, if you want to specify a record component C that shouldn't change, declaring it constant isn't allowed in Ada:

Listing 35:	constant	properties.ads
-------------	----------	----------------

```
package Constant Properties is
1
2
       type Rec is record
3
          C : constant Integer;
4
               ^^^^/
          - -
5
             ERROR: record components
          - - -
6
                      cannot be constant.
          - -
7
          V :
                        Integer:
8
      end record;
9
10
   end Constant Properties;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Constant_Properties
MD5: ba189f437348c5892847d067b0bc2e78

Build output

```
constant_properties.ads:4:11: error: constant component not permitted
gprbuild: *** compilation phase failed
```

A simple solution is to use a record discriminant:

Listing 36: constant properties.ads

```
1 package Constant_Properties is
2
3 type Rec (C : Integer) is
4 record
5 V : Integer;
6 end record;
7
8 end Constant_Properties;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Constant_Properties
MD5: b638c2fd78761def2b60e9ae7dceb765

A record discriminant can be accessed as a normal component, but it is read-only, so we cannot change it:

Listing 37: show_constant_property.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Constant Properties;
3
   use Constant_Properties;
4
5
   procedure Show Constant Property is
6
      R : Rec (10);
7
   begin
8
      Put_Line ("R.C = "
9
                 & R.C'Image);
10
11
      R.C := R.C + 1;
12
      -- ERROR: cannot change
13
                  record discriminant
14
  end Show_Constant_Property;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Constant_Properties
MD5: 82cde0032f2cb022e690f1175216fd77
```

Build output

```
show_constant_property.adb:12:05: error: assignment to discriminant not allowed
gprbuild: *** compilation phase failed
```

In this code example, the compilation fails because we cannot change the C discriminant. In this sense, C is a basically a constant component of the R object.

4.4.3 Private types

As we've seen in previous chapters, private types can have discriminants. For example:

Listing 38: private_with_discriminants.ads

```
package Private_With_Discriminants is
type T (L : Positive) is private;
```

```
private
5
6
       type Integer_Array is
7
         array (Positive range <>) of Integer;
8
9
      type T (L : Positive) is
10
       record
11
          Arr : Integer_Array (1 .. L);
12
       end record;
13
14
   end Private With Discriminants;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Private_With_

→Discriminants

MD5: 8f63443479e31a187a038381d9a32831
```

Here, discriminant L is used to specify the constraints of the array component Arr. Note that the same discriminant part must appear in both *the partial and the full view* (page 38) of type T.

4.4.4 Object declaration

As we've already seen, we declare objects of a type T with a discriminant D by specifying the actual value of discriminant D. This is called a *discriminant constraint* (page 211). For example:

Listing 39: recs.ads

```
package Recs is

type T (L : Positive;
 M : Positive) is
null record;

end Recs;
```

Listing 40: show_object_declaration.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Recs;
                       use Recs;
3
4
   procedure Show_Object_Declaration is
5
      A : T (L => 5, M => 6);
6
      B : T (7, 8);
7
   begin
8
      Put_Line ("A.L = "
9
                 & A.L'Image);
10
      Put_Line ("A.M = "
11
                 & A.M'Image);
12
      Put Line ("B.L = "
13
                 & B.L'Image);
14
      Put Line ("B.M = "
15
                 & B.M'Image);
16
   end Show_Object_Declaration;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Objects_ ⇔Discriminants MD5: 9daae29be9d0f99980ca152a3aca7363

Runtime output

A.L = 5 A.M = 6 B.L = 7 B.M = 8

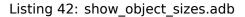
As we can see in the declaration of objects A and B, for the discriminant values, we can use a positional ((7, 8)) or named association $((L \Rightarrow 5, M \Rightarrow 6))$.

Object size

Discriminants can have an impact on the object size because we can set the discriminant to constraint a component of an *indefinite subtype* (page 31). For example:

Listing 41: recs.ads

```
package Recs is
1
2
       type Null_Rec (L : Positive;
3
                       M : Positive) is
4
         private;
5
6
       type Rec_Array (L : Positive) is
7
         private;
8
9
   private
10
11
       type Null_Rec (L : Positive;
12
                       M : Positive) is
13
         null record;
14
15
       type Integer_Array is
16
         array (Positive range <>) of Integer;
17
18
       type Rec_Array (L : Positive) is
19
       record
20
          Arr : Integer_Array (1 .. L);
21
       end record;
22
23
   end Recs;
24
```



```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Recs;
                      use Recs;
3
4
   procedure Show Object Sizes is
5
      Null Rec_A : Null_Rec (1, 2);
6
      Null Rec B : Null Rec (5, 6);
7
      Rec Array A : Rec Array (10);
8
      Rec_Array_B : Rec_Array (20);
9
   begin
10
      Put_Line ("Null_Rec_A'Size = "
11
                 & Null_Rec_A'Size'Image);
12
      Put_Line ("Null_Rec_B'Size = "
13
```

```
14 & Null_Rec_B'Size'Image);
15 Put_Line ("Rec_Array_A'Size = "
16 & Rec_Array_A'Size'Image);
17 Put_Line ("Rec_Array_B'Size = "
18 & Rec_Array_B'Size'Image);
19 end Show_Object_Sizes;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Objects_

⊲Discriminants_Size

MD5: 0abbc12286aff9fe428ea585564cf6d4
```

Build output

Runtime output

```
Null_Rec_A'Size = 64
Null_Rec_B'Size = 64
Rec_Array_A'Size = 352
Rec_Array_B'Size = 672
```

In this example, Null_Rec_A and Null_Rec_B have the same size because the type is a null record. However, Rec_Array_A and Rec_Array_B have different sizes because we're setting the L discriminant — which we use to constraint the Arr array component of the Rec_Array type — to 10 and 20, respectively.

4.4.5 Object assignments

As we've just seen, when we set the values for the discriminants of a type in the object declaration, we're constraining the objects. Those constraints are checked at runtime by the *discriminant check* (page 515). If the discriminants don't match, the Constraint_Error exception is raised.

Let's see an example:

Listing 43: recs.ads

```
1 package Recs is
2
3 type T (L : Positive;
4 M : Positive) is
5 null record;
6
7 end Recs;
```

Listing 44: show object assignments.adb

```
1 with Recs; use Recs;
2
3 procedure Show_Object_Assignments is
4 A1, A2 : T (5, 6);
5 B : T (7, 8);
6 begin
7 A1 := A2; -- OK
```

8 B := A1; -- ERROR!
9 end Show_Object_Assignments;

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Object_Assignments
MD5: 199247f1c0575c6845d85fd1911e1cf2

Build output

```
show_object_assignments.adb:8:10: warning: incorrect value for discriminant "L"_

→[enabled by default]

show_object_assignments.adb:8:10: warning: Constraint_Error will be raised at run_

→time [enabled by default]
```

Runtime output

raised CONSTRAINT_ERROR : show_object_assignments.adb:8 discriminant check failed

In this example, the A1 := A2 assignment is accepted because both A1 and A2 have the same constraints ((5, 6)). However, the B := A1 assignment is not accepted because the discriminant check fails at runtime.

Note that the discriminant check is not performed when we use *mutable subtypes* (page 199) — we discuss this specific kind of subtypes later on.

4.4.6 Discriminant type

In a discriminant specification, the type of the discriminant can only be a discrete subtype or an *access type* (page 603). Other kinds of types — e.g. composite types such as record types — are illegal for discriminants. However, we can always use them indirectly by using access types. (We'll see an example later.)

In addition to that, we can also use a different kind of access types, namely *anonymous access-to-object subtypes* (page 715). This specific kind of discriminant is called *access discriminant* (page 725). We discuss this topic in more details in another chapter.

Let's see a code example:

Listing 45: recs.ads

```
package Recs is
1
2
      type Usage Mode is (Off,
3
                            Simple Usage,
4
                            Advanced Usage);
5
6
      type Priv Info is private;
7
8
      type Priv Info Access is access Priv Info;
9
10
      type Proc Access is
11
        access procedure (P : in out Priv Info);
12
13
      type Priv Rec (Last : Positive;
14
                       Usage : Usage Mode;
15
                       Info : Priv_Info_Access;
16
                       Proc : Proc Access) is
17
         private;
18
19
```

```
private
20
21
       type Priv_Info is record
22
         A : Positive;
23
         B : Positive;
24
       end record;
25
26
       type Priv_Rec (Last : Positive;
27
                       Usage : Usage_Mode;
28
                       Info : Priv_Info_Access;
29
                       Proc : Proc Access) is
30
         null record;
31
32
   end Recs;
33
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Discriminants_

→Subtype

MD5: 4ddbc703d8ffcd6dc31e3715df62931a
```

In this example, we're declaring the Priv_Rec type with the following discriminants:

- The Last discriminant of the scalar (i.e. discrete) type Positive;
- The Usage discriminant of the enumeration (i.e. discrete) type Usage_Mode;
- The Info discriminant of the access-to-object type (page 593) Priv_Info_Access;
 - We discuss *access-to-object types as discriminant type* (page 603) in another chapter.
- The Proc discriminant of the access-to-subprogram type (page 677) Proc_Access;
 - We discuss access-to-subprogram types as discriminant type (page 683) in another chapter.

As indicated previously, it's illegal to use a private type or a record type as the type of a discriminant. For example:

Listing 46: recs.ads

```
package Recs is
1
2
      type Priv_Info is private;
3
4
5
      type Priv Rec (Info : Priv Info) is
        private;
6
                      ^^^^
7
          ERROR: cannot use private type
8
       - -
                  in discriminant.
       - -
9
10
   private
11
12
      type Priv_Info is record
13
        A : Positive;
14
        B : Positive;
15
      end record;
16
17
      type Priv_Rec (Info : Priv_Info) is
18
19
        null record;
20
   end Recs;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Discriminants_

→Subtype_Error

MD5: 17f36f0e09cb069d8215b38adbb46541
```

Build output

```
recs.ads:5:26: error: discriminants must have a discrete or access type gprbuild: *** compilation phase failed
```

We cannot use the Priv_Info directly as a discriminant type because it's a private type. However, as we've just seen in the previous code example, we use it indirectly by using an access type to this private type (see Priv_Info_Access in the code example).

Indefinite subtypes as discriminants

As we already implied, we cannot use indefinite subtypes as discriminants. For example, the following code won't compile:

Listing 47: unconstrained_types.ads

```
package Unconstrained Types is
1
2
      type Integer Array is
3
        array (Positive range <>) of Integer;
4
5
      type Simple_Record (Arr : Integer_Array) is
6
7
      -- ERROR: cannot use indefinite type
8
                 in discriminant.
9
      record
10
         L : Natural := Arr'Length;
11
      end record;
12
13
   end Unconstrained Types;
14
```

Code block metadata

Build output

```
unconstrained_types.ads:6:30: error: discriminants must have a discrete or access

→type

gprbuild: *** compilation phase failed
```

Integer_Array is a correct type declaration — although the type itself is indefinite after the declaration. However, we cannot use it as the discriminant in the declaration of Simple_Record. We could, however, have a correct declaration by using discriminants as access values:

Listing 48: unconstrained_types.ads

```
1 package Unconstrained_Types is
2
3 type Integer_Array is
4 array (Positive range <>) of Integer;
5
```

```
type Integer_Array_Access is
6
         access Integer_Array;
7
8
      type Simple_Record
9
         (Arr : Integer_Array_Access) is
10
      record
11
         L : Natural := Arr'Length;
12
      end record;
13
14
   end Unconstrained_Types;
15
```

Code block metadata

By adding the Integer_Array_Access type and using it in Simple_Record's type declaration, we can indirectly use an indefinite type in the declaration of another indefinite type. We discuss this topic later *in another chapter* (page 603).

4.4.7 Default values

We can specify default values for discriminants. Note, however, that we must either specify default values for **all** discriminants of the discriminant part or for none of them. This contrasts with default values for subprogram parameters, where we can *specify default values for just a subset of all parameters of a specific subprogram* (page 474).

As expected, we can override the default values by specifying the values of each discriminant when declaring an object. Let's see a simple example:

Listing 49: recs.ads

```
package Recs is
1
2
       type T (L : Positive := 1;
3
               M : Positive := 2) is
4
         private;
5
6
   private
7
8
       type T (L : Positive := 1;
9
               M : Positive := 2) is
10
         null record;
11
12
   end Recs;
13
```

Listing 50: show_object_declaration.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Recs;
                       use Recs;
3
4
   procedure Show Object Declaration is
5
      A : T;
6
      B : T (7, 8);
7
   begin
8
      Put_Line ("A.L = "
9
                 & A.L'Image);
10
      Put_Line ("A.M = "
11
```

```
12 & & A.M'Image);
13 Put_Line ("B.L = "
14 & & B.L'Image);
15 Put_Line ("B.M = "
16 & & B.M'Image);
17 end Show_Object_Declaration;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Discriminant_

→Default_Value

MD5: 33385c4ba4ed9fc90c55990bde0b70cb
```

Runtime output

A.L = 1 A.M = 2 B.L = 7 B.M = 8

In this example, object A makes use of the default values for the discriminants of type T, so it has the discriminants (L => 1, M => 2). In the case of object B, we're specifying the values (L => 7, M => 8), which are used instead of the default values.

Note that we cannot set default values for nonlimited tagged types. The same applies to generic formal types. For example:

```
Listing 51: recs.ads
```

```
package Recs is
1
2
      type TT (L : Positive := 1;
3
                M : Positive := 2) is
4
                ~~~~~~~~~~~
5
          ERROR: cannot assign default
6
                  in discriminant of
       - -
7
                  nonlimited tagged type.
8
        tagged private;
9
10
      type LTT (L : Positive := 1;
11
                 M : Positive := 2) is
12
         tagged limited private;
13
14
   private
15
16
      type TT (L : Positive := 1;
17
                M : Positive := 2) is
18
         tagged null record;
19
20
      type LTT (L : Positive := 1;
21
                 M : Positive := 2) is
22
         tagged limited null record;
23
24
   end Recs;
25
```

Code block metadata

Build output

recs.ads:3:29: error: discriminants of nonlimited tagged type cannot have defaults recs.ads:4:29: error: discriminants of nonlimited tagged type cannot have defaults gprbuild: *** compilation phase failed

As we can see, compilation fails because of the default values for the discriminants of the nonlimited tagged type TT. In the case of the limited tagged type LTT, the default values for the discriminants are legal.

Mutable subtypes

An unconstrained discriminated subtype with defaults is called a mutable subtype, and a variable of such a subtype is called a mutable variable because the discriminants of such a variable can be changed. An important feature of mutable subtypes is that it allows for changing the discriminants of an object via assignments — in this case, no *discriminant check* (page 515) is performed.

Let's see an example:

Listing 52: mutability.ads

```
package Mutability is
1
2
       type T_Non_Mutable
3
         (L : Positive;
4
         M : Positive) is
5
         null record;
6
7
      type T Mutable
8
         (L : Positive := 1;
9
         M : Positive := 2) is
10
         null record;
11
12
   end Mutability;
13
```

Listing 53: show_mutable_subtype_assignment.adb

```
with Mutability; use Mutability;
1
2
   procedure Show_Mutable_Subtype_Assignment is
3
      NM_1 : T_Non_Mutable (5, 6);
4
      NM_2 : T_Non_Mutable (7, 8);
5
6
      M 1 : T Mutable (7, 8);
7
      M_2 : T_Mutable;
8
   begin
9
      NM_2 := NM_1; -- ERROR!
10
      M \overline{2} := M \overline{1};
                       -- OK
11
   end Show Mutable Subtype Assignment;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Mutable_Subtype
MD5: ace4470544bdc6efb0dca7039ca33cbc

Build output

Runtime output

In this example, the NM_2 := NM_1 assignment fails because both objects are of a nonmutable subtype with different discriminants, so that the discriminant check fails at runtime. However, the M_2 := M_1 assignment is OK because both objects are mutable variables. In this case, this assignment changes the discriminants of M_2 from (L => 1, M => 2) to (L => 7, M => 8).

Note that assignments of mutable variables may not always work at runtime. For example, if a discriminant of a mutable subtype is used to constraint a component of indefinite subtype, we might see the corresponding checks fail at runtime. For example:

Listing 54: mutability.ads

```
package Mutability is
1
2
       type T Mutable Array (L : Positive := 10) is
3
         private;
4
5
   private
6
7
      type Integer_Array is
8
         array (Positive range <>) of Integer;
9
10
      type T_Mutable_Array (L : Positive := 10) is
11
       record
12
          Arr : Integer Array (1 .. L);
13
      end record;
14
15
   end Mutability;
16
```

Listing 55: show_mutable_subtype_error.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Mutability; use Mutability;
3
4
   procedure Show_Mutable_Subtype_Error is
5
      A : T_Mutable_Array (10);
6
      B : T_Mutable_Array (20);
7
   begin
8
      Put_Line ("A'Size = "
9
                 & A'Size'Image);
10
      Put_Line ("B'Size = "
11
                 & B'Size'Image);
12
13
      A := B; -- ERROR!
14
   end Show_Mutable_Subtype_Error;
15
```

Code block metadata

Build output

Runtime output

A'Size = 352 B'Size = 672

raised CONSTRAINT_ERROR : show_mutable_subtype_error.adb:14 discriminant check_ ⊣failed

In this case, the assignment A := B raises the Constraint_Error exception at runtime. Here, the Arr component of each object has a different range: 1 . . 10 for object A and 1 . . 20 for object B. To prevent this situation, we should declare T_Mutable_Array as a limited type, so that assignments are not permitted.

4.4.8 Derived types and subtypes

As expected, we may derive types with discriminants or declare subtypes of it. However, there are a couple of details associated with this, which we discuss now.

Subtypes

When declaring a subtype of a type with discriminants, we have the choice to specify the value of the discriminants for the parent type, or specify no discriminants at all:

Listing 56: subtypes_with_discriminants.ads

```
package Subtypes With Discriminants is
1
2
       type T
3
         (L : Positive;
4
         M : Positive) is
5
         null record;
6
7
       subtype Sub_T is T;
8
       -- Discriminants are not specified:
9
       -- taking the ones from T.
10
11
       subtype Sub T 2 is T
12
       (L \Rightarrow 3, M \Rightarrow 4);
13
       -- Discriminants are specified:
14
          taking the ones from Sub_T_2
       - -
15
16
   end Subtypes_With_Discriminants;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Subtypes
MD5: 6f02c295f295c81fd20d06f7c710994c

For the Sub_T subtype declaration in this example, we don't specify values for the parent type's discriminants. For Sub_T_2, in contrast, we set the discriminants to (L => 3, M => 4).

When declaring objects of these subtypes, we need to take the constraints into account:

```
Listing 57: subtypes_with_discriminants.ads
```

```
package Subtypes With Discriminants is
1
2
       type T
3
         (L : Positive;
4
         M : Positive) is
5
         null record;
6
7
       subtype Sub_T is T;
8
       -- Discriminants are not specified:
9
          taking the ones from T.
       - -
10
11
       subtype Sub_T_2 is T
12
       (L \Rightarrow 3, M \Rightarrow 4);
13
       -- Discriminants are specified:
14
       -- taking the ones from Sub_T_2
15
16
   end Subtypes_With_Discriminants;
17
```

Listing 58: show_subtypes_with_discriminants.adb

```
with Subtypes_With_Discriminants;
1
   use Subtypes_With_Discriminants;
2
3
   procedure Show_Subtypes_With_Discriminants is
4
      A1 : T (1, 2);
5
      A2 : T (3, 4);
6
      B1 : Sub_T (1, 2);
7
      B2 : Sub_T (3, 4);
8
      C2 : Sub_T_2;
9
10
       -- C1 : Sub_T_2 (1, 2);
11
       - -
12
       -- ERROR: discriminants already
13
       - -
                  constrained
14
   begin
15
      B1 := A1;
16
      -- OK: discriminants match
17
18
      B2 := A1;
19
      -- CONSTRAINT_ERROR!
20
21
      B2 := A2;
22
      -- OK: discriminants match
23
24
      C2 := A1;
25
       -- CONSTRAINT_ERROR!
26
27
      C2 := A2;
28
       -- OK: discriminants match
29
   end Show_Subtypes_With_Discriminants;
30
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Subtypes
MD5: 9a8516d70e7a53ae332e5c5b6df7f04e

Build output

```
discriminant "L" [enabled by default]
show_subtypes_with_discriminants.adb:19:10: warning: Constraint_Error will be_
oraised at run time [enabled by default]
show_subtypes_with_discriminants.adb:25:10: warning: incorrect value for_
odiscriminant "L" [enabled by default]
show_subtypes_with_discriminants.adb:25:10: warning: Constraint_Error will be_
oraised at run time [enabled by default]
```

Runtime output

```
raised CONSTRAINT_ERROR : show_subtypes_with_discriminants.adb:19 discriminant

⇔check failed
```

For objects of Sub_T subtype, we *have to* specify the value of each discriminant. On the other hand, for objects of Sub_T_2 type, we *cannot* specify the constraints because they have already been defined in the subtype's declaration — in this case, they're always set to (3, 4).

When assigning objects of different subtypes, the discriminant check will be performed — as we *mentioned before* (page 193). In this example, the assignments B2 := A1 and C2 := A1 fail because the objects have different constraints.

Derived types

The behavior for derived types is very similar to the one we've just described for subtypes. For example:

Listing 59: derived_with_discriminants.ads

```
package Derived With Discriminants is
1
2
       type T
3
         (L : Positive;
4
          M : Positive) is
5
         null record;
6
7
       type T_Derived is new T;
8
       -- Discriminants are not specified:
9
       -- taking the ones from T.
10
11
       type T_Derived_2 is new T
12
        (L \Rightarrow 3, M \Rightarrow 4);
13
       -- Discriminants are specified:
14
          taking the ones from T Derived 2
15
16
17
```

end Derived_With_Discriminants;

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types
MD5: 1e88f787bd9b568e43fc423c121f24f7

For the T_Derived type, we reuse the discriminants of the parent type T. For the T_Derived_2 type, we specify a value for each discriminant of T.

As you probably notice, this code looks very similar to the code using subtypes. The main difference between using subtypes and derived types is that, as expected, we have to perform a *type conversion* (page 45) in the assignments:

```
Listing 60: show_derived_with_discriminants.adb
```

```
with Derived With Discriminants;
1
   use Derived_With_Discriminants;
2
3
   procedure Show_Derived_With_Discriminants is
4
      A1 : T (1, 2);
5
      A2 : T (3, 4);
6
      B1 : T_Derived (1, 2);
7
      B2 : T_Derived (3, 4);
8
      C2 : T Derived 2;
9
10
       -- C1 : Sub_T_2 (1, 2);
11
       - -
12
       -- ERROR: discriminants already
13
       - -
                  constrained
14
   begin
15
      B1 := T_Derived (A1);
16
      -- OK: discriminants match
17
18
      B2 := T Derived (A1);
19
      -- ERROR!
20
21
      C2 := T Derived 2 (A1);
22
      -- CONSTRAINT_ERROR!
23
24
      C2 := T_Derived_2 (A2);
25
      -- OK: discriminants match
26
  end Show Derived With Discriminants;
27
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types
MD5: 1b32807fcd3b343fbf8ab0d0287ca5bb

Build output

Runtime output

```
raised CONSTRAINT_ERROR : show_derived_with_discriminants.adb:19 discriminant

→check failed
```

Once again, a discriminant check is performed when assigning objects to ensure that the type discriminants match. In this code example, the assignments B2 := A1 and C2 := A1 fail because the objects have different constraints.

Derived types with renamed discriminants

We could rewrite a type declaration such as **type T_Derived is new** T by explicitly declaring the discriminants. We can do that for the previous code example:

Listing 61: derived_with_discriminants.ads

```
package Derived_With_Discriminants is
```

2

```
type T
3
          (L : Positive;
4
          M : Positive) is
5
         null record;
6
7
           The declaration:
8
       - -
9
       - -
                type T Derived is new T;
       - -
10
        - -
11
           is the same as:
       - -
12
13
       type T_Derived
14
15
          (L : Positive;
          M : Positive) is
16
          new T (L \Rightarrow L, M \Rightarrow M);
17
18
   end Derived_With_Discriminants;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_Same_

⇔Discriminants

MD5: 3ee4b3a70e8ab9ba2684c6a2c695f689
```

We may, however, rename the discriminants instead. For example, we could rename L and M to X and Y. For example:

Listing 62: derived_with_discriminants.ads

```
package Derived With Discriminants is
1
2
3
       type T
         (L : Positive;
4
          M : Positive) is
5
         null record;
6
7
       type T Derived
8
         (X : Positive;
9
          Y : Positive) is
10
         new T (L \Rightarrow X, M \Rightarrow Y);
11
12
   end Derived With Discriminants;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_

→Renamed_Discriminants

MD5: ec2954f538fa63b4d3c7c134527be35d
```

Of course, if we use named association when declaring objects, we have to use the correct discriminant names:

Listing 63: show_derived_with_discriminants.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Derived_With_Discriminants;
use Derived_With_Discriminants;

procedure Show_Derived_With_Discriminants is
A : T (L => 1, M => 2);
```

```
B : T_Derived (X => 3, Y => 4);
8
                       ~~~~~~
9
      -- Using correct discriminant names
10
   begin
11
      Put_Line ("A.L = "
12
                 & A.L'Image);
13
      Put_Line ("A.M = "
14
                 & A.M'Image);
15
      Put_Line ("B.X = "
16
                 & B.X'Image);
17
      Put Line ("B.Y = "
18
                 & B.Y'Image);
19
20
   end Show_Derived_With_Discriminants;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_ ⊲Renamed_Discriminants MD5: cd3ca2d84c8b7d334b152ebab1955a5e

Runtime output

A.L = 1 A.M = 2 B.X = 3 B.Y = 4

In essence, the discriminants of both parent and derived types are the same: the only difference is that they are accessed by different names. This allows us to convert from a parent type to a derived type:

Listing 64: show_derived_with_discriminants.adb

```
with Derived With Discriminants;
1
   use Derived With Discriminants;
2
3
   procedure Show Derived With Discriminants is
4
      A : T (L \implies 1, M \implies 2);
5
      B : T Derived (X \Rightarrow 1, Y \Rightarrow 2);
6
   begin
7
      B := T_Derived (A); -- OK
8
   end Show Derived With Discriminants;
a
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_

→Renamed_Discriminants

MD5: d685f16bf3a9d64b4c1f182880455ad0
```

Here, even though objects A and B have discriminants with different names, the assignment $B := T_{Derived}$ (A) is valid.

Derived types with more constrained discriminants

When deriving types with discriminants, we may use a more constrained type for the discriminants of derived type. For example, if the discriminant D of the parent type is of **Integer** type, the corresponding discriminant of the derived type may use a constrained subtype such as **Natural** or **Positive** — because both **Natural** and **Positive** are subtypes of type **Integer**. For example: Listing 65: derived_with_discriminants.ads

```
package Derived With Discriminants is
1
2
       type T
3
          (L : Integer;
4
          M : Integer) is
5
         null record;
6
7
       type T_Derived_2
8
          (X : Natural;
9
          Y : Positive) is
10
          new T (L \Rightarrow X, M \Rightarrow Y);
11
12
   end Derived_With_Discriminants;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_More_

⊶Constrained_Discriminants

MD5: 413f04f1f98dde2a6e0df3ee6955da7f
```

As expected, the constraints of each discriminant's type are taken into account when evaluating the value that is specified for each discriminant:

Listing 66: show_derived_with_discriminants.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Derived_With_Discriminants;
3
   use Derived_With_Discriminants;
4
5
   procedure Show_Derived_With_Discriminants is
6
      A : T (L => -1, M => -2);
7
      B : T_Derived_2 (X => 0, Y => 1);
8
   begin
9
      Put_Line ("A.L = "
10
                 & A.L'Image);
11
      Put Line ("A.M = "
12
                 & A.M'Image);
13
      Put_Line ("B.X = "
14
                 & B.X'Image);
15
      Put_Line ("B.Y = "
16
                 & B.Y'Image);
17
   end Show_Derived_With_Discriminants;
18
```

Code block metadata

Runtime output

A.L = -1 A.M = -2 B.X = 0B.Y = 1

Here, we can use $(L \Rightarrow -1, M \Rightarrow -2)$ in the declaration of object A because both discriminants are of **Integer** type. However, in the declaration of object B, we can only use values for the discriminants that are in the range of the **Natural** and **Positive** subtypes,

respectively. (If you change the code to use negative values instead, a Constraint_Error exception is raised at runtime.)

Extending the discriminant part

As we've seen, we can rename discriminants or use more constrained subtypes for discriminants in derived types. We might also want to add a new discriminant to the derived type — in addition to the discriminants of the parent's type. However, this is considered a type extension, as the new discriminant is part of the type definition.

As an example, we may want to add the A discriminant of **Boolean** type to a derived type. For non-tagged types, such a declaration will trigger a compilation error as expected:

Listing 67:	derived	with	discriminants.ads

```
package Derived With Discriminants is
1
2
       type T
3
         (L : Positive;
4
         M : Positive) is
5
         null record;
6
7
       type T Derived
8
         (X : Positive;
9
          Y : Positive;
10
         A : Boolean) is
11
           ~~~~~
12
       -- ERROR: cannot extend type with new
13
                  Boolean discriminant A
       - -
14
         new T (L \Rightarrow X, M \Rightarrow Y);
15
16
   end Derived With Discriminants;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_

→Extension_Error

MD5: f9ba4ae1c344d63ed706005e30fe60c2
```

Build output

```
derived_with_discriminants.ads:11:07: error: new discriminants must constrain old

ones

gprbuild: *** compilation phase failed
```

To circumvent this issue, we could, of course, declare a component of T type instead of deriving from it:

Listing 68: derived_with_discriminants.ads

```
package Derived With Discriminants is
1
2
       type T
3
         (L : Positive;
4
         M : Positive) is
5
         null record;
6
7
      type T 2
8
         (X : Positive;
9
          Y : Positive;
10
         A : Boolean) is
11
      record
12
```

```
13 A_Comp : T (L => X, M => Y);
14 end record;
15
16 end Derived_With_Discriminants;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_

→Extension_Error

MD5: 41e911cdc8486cd49931b2082586d8e7
```

In this case, A_Comp is a component of type T, and we're using the discriminant X and Y as the constraints of this component.

Naturally, using tagged types is another alternative:

```
Listing 69: derived_with_discriminants.ads
```

```
package Derived_With_Discriminants is
1
2
3
       type T
         (L : Positive;
4
          M : Positive) is
5
         tagged null record;
6
7
8
       type T_Derived_Extended
         (X : Positive;
9
          Y : Positive;
10
          A : Boolean) is -- New discriminant
11
         new T (L \Rightarrow X, M \Rightarrow Y)
12
           with null record;
13
14
       type T Derived Extended 2
15
         (A : Boolean;
                             -- New discriminant
16
          X : Positive;
17
          Y : Positive) is
18
         new T (L \Rightarrow X, M \Rightarrow Y)
19
           with null record;
20
21
       type T_Derived_Extended_3
22
         (A : Boolean) is -- New discriminant
23
         new T (L => 1, M => 2)
24
           with null record;
25
26
       type T_Derived_Extended_4
27
         (A : Boolean;
                             -- New discriminant
28
          X : Positive) is
29
         new T (L \Rightarrow X, M \Rightarrow X)
30
           with null record;
31
32
   end Derived_With_Discriminants;
33
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Tagged_Types
MD5: b8124d132a4b5066826980c8cc43a7ad

In this code example, we're adding the A discriminant when declaring T_Derived_Extended. Because T is a tagged type, such a new discriminant is fine.

Note that the order of the discriminants can be rearranged: when deriving a new type, we don't need to specify the discriminants of the parent type before any new discriminants. In

fact, in the declaration of T_Derived_Extended_2, the additional discriminant A is declared before the discriminants that match the parent type's discriminants.

In addition, we may even use literals to specify the constraints for the parent type — as we're doing in the declaration of T_Derived_Extended_3. Also, we can use the same discriminant from the derived type for the constraints of the parent type — in the declaration of T_Derived_Extended_4, we use the X discriminant for both L and M discriminants of type T.

Deriving with defaults

If the discriminants of the parent type have default values, those default values are inherited by the derived type. Alternatively, we can set different default values.

Let's see a code example:

```
Listing 70: derived_with_discriminants.ads
```

```
package Derived With Discriminants is
1
2
       type T
3
         (L : Positive := 1;
4
          M : Positive := 2) is
5
         null record;
6
7
       type T Derived is new T;
8
9
       type T_Derived_2
10
         (L : Positive := 1;
11
          M : Positive := 3) is
12
         new T (L \Rightarrow L, M \Rightarrow M);
13
14
   end Derived_With_Discriminants;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_

⇔Defaults

MD5: 4ffa513c0cd8b6812359a2fc4d8325d2
```

In this example, the derived type T_Derived has the same default values as the parent type T, namely (L => 1, M => 2). For the derived type T_Derived_2, we're changing the value of M to 3 and keeping the same value for L.

As we've seen before, instead of setting default values, we can set the constraints of the parent type in the declaration of the derived type:

```
Listing 71: derived_with_discriminants.ads
```

```
package Derived_With_Discriminants is
1
2
       type T
3
         (L : Positive := 1;
4
          M : Positive := 2) is
5
         null record;
6
7
       type T_Derived_Constrainted is new T
8
         (L \implies 1, M \implies 3);
9
10
  end Derived_With_Discriminants;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_ →Defaults_Constraints MD5: d6053fc79a3e7010ec7b3ec73f51f4e5

In this case, we're constraining the discriminants of the parent type to $(L \Rightarrow 1, M \Rightarrow 3)$. Note that L has the same value as the default value set for the parent type T.

For further reading...

In other contexts (such as *record aggregates* (page 249), which we discuss in another chapter), we could use the so-called *box notation* (page 252) to specify that we want to use the default value. This, however, isn't possible with type discriminants:

```
Listing 72: derived with discriminants.ads
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants.Derived_Types_

→Defaults_Constraints_Box_Notation

MD5: 18d755ec4de45164a47009ab25368452
```

Build output

1 2

3

4

5

6 7

8

9 10

11

12

13

```
derived_with_discriminants.ads:9:11: error: missing operand
gprbuild: *** compilation phase failed
```

Instead of using <>, we have to repeat the value explicitly.

4.5 Discriminant constraints and operations

In this section, we discuss some details about discriminant constraints and operations related to discriminants — more specifically, the Constrained attribute.

In the Ada Reference Manual

• 3.7.1 Discriminant Constraints⁹⁷

4.5.1 Discriminant constraints

As we discussed before, when *declaring an object with a discriminant* (page 191), we have to specify the values of the all discriminants — unless, of course, those discriminants have a *default value* (page 197). The values we specify for the discriminants are called discriminant constraints.

```
<sup>97</sup> http://www.ada-auth.org/standards/12rm/html/RM-3-7-1.html
```

Let's revisit the code example we've seen earlier on:

```
Listing 73: recs.ads
```

```
1 package Recs is
2
3 type T (L : Positive;
4 M : Positive) is
5 null record;
6
7 end Recs;
```

Listing 74: show_object_declaration.adb

```
with Recs;
                      use Recs;
1
2
  procedure Show_Object_Declaration is
3
      A : T (L => 5, M => 6);
4
      B : T (7, 8);
5
      C : T (7, M => 8);
6
  begin
7
      null;
8
  end Show_Object_Declaration;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants_Constraints_

⊲Operations.Discriminant_Constraint

MD5: 9e37a1cde73f27b99fd2a9eb57f23c44
```

Here, L => 5, M => 6 (for object A) are named constraints, while 7, 8 (for object B) are positional constraints.

It's possible to use both positional and named constraints, as we do for object C: 7, $M \Rightarrow 8$. In this case, the positional associations must precede the named associations.

In the case of named constraints, we can use multiple selector names:

Listing 75: show_object_declaration.adb

```
with Recs;
                     use Recs;
1
2
  procedure Show Object Declaration is
3
      A : T (L | M => 5);
4
5
      -- multiple selector names
6
  begin
7
      null;
8
  end Show_Object_Declaration;
a
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants_Constraints_

Goperations.Discriminant_Constraint

MD5: b6fbe1d69bb520a7b6845536a1601978
```

This is only possible, however, if those named discriminants are all of the same type. (In this case, L and M are both of **Positive** subtype.)

1 In the Ada Reference Manual

• 3.7.1 Discriminant Constraints⁹⁸

Discriminant constraint in subtypes

We can use discriminant constraints in the declaration of subtypes. For example:

```
Listing 76: show object declaration.adb
```

```
with Recs;
                      use Recs;
1
2
   procedure Show Object Declaration is
3
      subtype T_5_6 is T (L => 5, M => 6);
4
                           ~~~~~~
5
      -- discriminant constraints for subtype
6
7
      A : T_5_6;
8
   begin
9
      null:
10
   end Show Object Declaration;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants_Constraints_

⊲Operations.Discriminant_Constraint

MD5: c6c4226073f4282d68d621519ca4d420
```

In this example, we use the named discriminant constraints L => 5, M => 6 in the declaration of the subtype T_5_6.

4.5.2 Constrained Attribute

We can use the Constrained attribute to verify whether an object of discriminated type is constrained or not. Let's look at a simple example:

Listing 77: recs.ads

```
1 package Recs is
2
3 type T (L : Positive := 1) is
4 null record;
5
6 end Recs;
```

Listing 78: show_constrained_attribute.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Recs;
                      use Recs;
3
4
   procedure Show_Constrained_Attribute is
5
      Constr : T (L => 5);
6
                      ^^^^ constrained.
7
      Unconstr : T;
8
                  ^ unconstrained;
      - -
9
      - -
                   using defaults.
10
   begin
11
      Put_Line ("Constr'Constrained:
                                          н
12
                 & Constr'Constrained'Image);
13
      Put_Line ("Unconstr'Constrained: "
14
                 & Unconstr'Constrained'Image);
15
   end Show_Constrained_Attribute;
16
```

Code block metadata

98 http://www.ada-auth.org/standards/22rm/html/RM-3-7-1.html

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants_Constraints_ Goperations.Simple_Constrained_Attribute MD5: 6a9a807f5af132a07949d2887fa5bfe5

Runtime output

Constr'Constrained: TRUE Unconstr'Constrained: FALSE

As the Constrained attribute indicates, the Constr object is constrained (by the L => 5 discriminant constraint), while the Unconstr object is unconstrained. Note that, even though Unconstr is using the default value for L — which would correspond to the discriminant constraint L => 1 — the object itself hasn't been constraint at its declaration.

Let's continue our discussion with a more complex example by reusing the Unconstrained_Types package that we declared in a *previous section* (page 31). In this version of the package, we're adding a Reset procedure for the discriminated record type Simple_Record:

Listing 79: unconstrained_types.ads

```
package Unconstrained_Types is
1
2
       type Simple Record
3
         (Extended : Boolean := False) is
4
       record
5
          V : Integer;
6
          case Extended is
7
             when False =>
8
                null;
9
             when True =>
10
                V Float : Float;
11
          end case;
12
       end record;
13
14
       procedure Reset (R : in out Simple_Record);
15
16
   end Unconstrained_Types;
17
```

Listing 80: unconstrained_types.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Unconstrained_Types is
3
4
      procedure Reset (R : in out Simple_Record) is
5
          Zero Not_Extended : constant
6
            Simple_Record := (Extended => False,
7
                               V
                                         => 0):
8
9
          Zero_Extended : constant
10
            Simple_Record := (Extended => True,
11
                               v
                                       => 0,
12
                               V_Float => 0.0);
13
      begin
14
          Put Line ("---- Reset: R'Constrained => "
15
                    & R'Constrained'Image);
16
17
          if not R'Constrained then
18
             R := Zero_Extended;
19
          else
20
```

```
if R.Extended then
21
                 R := Zero_Extended;
22
              else
23
                 R := Zero_Not_Extended;
24
              end if;
25
          end if;
26
       end Reset;
27
28
   end Unconstrained_Types;
29
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants_Constraints_ Goperations.Constrained_Attribute MD5: b56e6d71fd4f05e8490412d7fe40b923

As the name indicates, the Reset procedure initializes all record components with zero. Note that we use the Constrained attribute to verify whether objects are constrained before assigning to them. For objects that are not constrained, we can simply assign another object to it — as we do with the R := Zero_Extended statement. When an object is constrained, however, the discriminants must match. If we assign an object to R, the discriminant of that object must match the discriminant of R. This is the kind of verification that we do in the **else** part of that procedure: we check the state of the Extended discriminant before assigning an object to the R parameter.

Note that the Simple_Record type has a *variant part* (page 227). We discuss this topic later on in this chapter.

Note as well that, in the initialization of the Zero_Not_Extended and Zero_Extended constants, we have to indicate the discriminant as a component of the aggregates (e.g.: (Extended => **False**, V => 0). We discuss this topic in another chapter when we learn more about *aggregates and record discriminants* (page 258).

The Using_Constrained_Attribute procedure below declares two objects of Simple_Record type: R1 and R2. Because the Simple_Record type has a default value for its discriminant, we can declare objects of this type without specifying a value for the discriminant. This is exactly what we do in the declaration of R1. Here, we don't specify any constraints, so that it takes the default value (Extended => **False**). In the declaration of R2, however, we explicitly set Extended to **False**:

Listing 81: using_constrained_attribute.adb

1	<pre>with Ada.Text_I0; use Ada.Text_I0;</pre>
2	
3	<pre>with Unconstrained_Types; use Unconstrained_Types;</pre>
4	
5	<pre>procedure Using_Constrained_Attribute is</pre>
6	R1 : Simple Record;
7	R2 : Simple_Record (Extended => False);
8	
9	procedure Show_Rs is
10	begin
11	<pre>Put_Line ("R1'Constrained => "</pre>
12	<pre>& R1'Constrained'Image);</pre>
13	<pre>Put_Line ("R1.Extended => "</pre>
14	<pre>& R1.Extended'Image);</pre>
15	<pre>Put_Line ("");</pre>
16	<pre>Put_Line ("R2'Constrained => "</pre>
17	& R2'Constrained'Image);
18	<pre>Put_Line ("R2.Extended => "</pre>
19	<pre>& R2.Extended'Image);</pre>

```
Put_Line ("-----");
20
      end Show_Rs;
21
   begin
22
      Show_Rs;
23
24
      Reset (R1);
25
      Reset (R2);
26
      Put_Line ("-----");
27
28
      Show Rs;
29
  end Using Constrained Attribute;
30
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Discriminants_Constraints_

Goperations.Constrained_Attribute

MD5: f7517fcd3c68a784f55064f188d4e7bb
```

Runtime output

```
R1'Constrained => FALSE
R1.Extended => FALSE
...
R2'Constrained => TRUE
R2.Extended => FALSE
.... Reset: R'Constrained => FALSE
.... Reset: R'Constrained => TRUE
....
R1'Constrained => FALSE
R1.Extended => TRUE
...
R2'Constrained => TRUE
...
```

When we run this code, the user messages from Show_Rs indicate to us that R1 is not constrained, while R2 is constrained. Because we declare R1 without specifying a value for the Extended discriminant, R1 is not constrained. In the declaration of R2, on the other hand, the explicit value for the Extended discriminant makes this object constrained. Note that, for both R1 and R2, the value of Extended is **False** in the declarations.

As we were just discussing, the Reset procedure includes checks to avoid mismatches in discriminants. When we don't have those checks, we might get exceptions at runtime. We can force this situation by replacing the implementation of the Reset procedure with the following lines:

```
-- [...]
begin
    Put_Line ("---- Reset: R'Constrained => "
        & R'Constrained'Image);
    R := Zero_Extended;
end Reset;
```

Running the code now generates a runtime exception:

raised CONSTRAINT_ERROR : unconstrained_types.adb:12 discriminant check failed

This exception is raised during the call to Reset (R2). As we see in the code, R2 is constrained. Also, its Extended discriminant is set to **False**, which means that it doesn't have the V_Float component. Therefore, R2 is not compatible with the constant Zero_Extended object, so we cannot assign Zero_Extended to R2. Also, because R2 is constrained, its Extended discriminant cannot be modified.

The behavior is different for the call to Reset (R1), which works fine. Here, when we pass R1 as an argument to the Reset procedure, its Extended discriminant is **False** by default. Thus, R1 is also not compatible with the Zero_Extended object. However, because R1 is not constrained, the assignment modifies R1 (by changing the value of the Extended discriminant). Therefore, with the call to Reset, the Extended discriminant of R1 changes from **False** to **True**.

```
In the Ada Reference Manual
```

3.7.2 Operations of Discriminated Types⁹⁹

4.6 Unknown discriminants

As we've seen *previously* (page 188), a type with discriminants can have known discriminants or unknown discriminants. In this section, we focus on unknown discriminants. Because the discriminants are unknown, this is an *indefinite type* (page 31). Let's start with a simple example:

Listing 82: unknown_discriminants.ads

```
package Unknown Discriminants is
1
2
      type T_Unknown_Discr (<>) is
3
4
          Unknown discriminant part
5
         private;
6
7
   private
8
9
      type T Unknown Discr is
10
        null record;
11
12
   end Unknown_Discriminants;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Simple_ →Example MD5: 5f673c957132b1bca633c247f857e37b

Note that we can only use an unknown discriminant part in the *partial view* (page 38); we cannot use it in the full view of a type:

Listing 83: unknown discriminants.ads

```
1 package Unknown_Discriminants is
2
3 type T_Unknown_Discr (<>) is
4 null record;
5
6 end Unknown Discriminants;
```

Code block metadata

99 http://www.ada-auth.org/standards/22rm/html/RM-3-7-2.html

Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Wrong_Full_ ⇔View MD5: dfce1471556af87b6a99314b1ee32446

Build output

```
unknown_discriminants.ads:3:25: error: full type declaration cannot have unknown

discriminants

gprbuild: *** compilation phase failed
```

To be more precise, an unknown discriminant part can only be used in the declaration of a private type, a private extension or an *incomplete type* (page 36). In addition, as we'll see in another chapter, it can also be used in the generic equivalents: generic private types, generic private extensions, generic incomplete types, and formal derived types.

For example:

Listing 84: unknown discriminants.ads	Listing 84:	unknown	discriminants.ads
---------------------------------------	-------------	---------	-------------------

```
package Unknown Discriminants is
1
2
           Private type
3
      type Rec (<>) is
4
5
         private;
6
            Tagged private type
7
      type Tagged_Rec (<>) is
8
        tagged private;
9
10
            Incomplete type
11
      type T Incomplete (<>);
12
13
       type T Incomplete (<>) is
14
         private;
15
16
   private
17
18
      type Rec is
19
         null record;
20
21
      type Tagged Rec is
22
         tagged null record;
23
24
       type T Incomplete is
25
         null record;
26
27
   end Unknown Discriminants;
28
```

Code block metadata

In this example, we have three forms of private types using an unknown discriminant part: an untagged private type (Rec), a tagged type (Tagged_Rec) and an incomplete type (T_Incomplete) that becomes an untagged private type.

1 In the Ada Reference Manual

• 3.7 Discriminants¹⁰⁰

4.6.1 Object declaration

Now, let's talk about objects of types with unknown discriminants. Consider the Rec type below:

```
Listing 85: unknown_discriminants.ads
```

```
package Unknown Discriminants is
1
2
      type Rec (<>) is private;
3
4
   private
5
6
      type Rec is
7
       record
8
          I : Integer;
9
      end record;
10
11
   end Unknown Discriminants;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Object_

→Declaration

MD5: 9f588870ec70ea30c795a6a0a602f589
```

We cannot declare objects of type Rec *directly*, as this type is *indefinite* (page 31):

Listing 86: show_object_declaration.adb

```
1 with Unknown_Discriminants;
2 use Unknown_Discriminants;
3 
4 procedure Show_Object_Declaration is
5 A : Rec;
6 begin
7 null;
8 end Show_Object_Declaration;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Object_

→Declaration

MD5: 5f30773fc17096943939468faf50338b
```

Build output

```
show_object_declaration.adb:5:08: error: unconstrained subtype not allowed (need_
__initialization)
gprbuild: *** compilation phase failed
```

Because the type is indefinite, it requires explicit initialization — we can do this by introducing a subprogram that initializes the type. In our code example, we can implement a simple Init function for this type:

Listing 87: unknown_discriminants.ads

```
1 package Unknown_Discriminants is
2
3 type Rec (<>) is private;
4
```

(continues on next page)

¹⁰⁰ http://www.ada-auth.org/standards/12rm/html/RM-3-7.html

```
function Init return Rec;
5
6
   private
7
8
       type Rec is
9
       record
10
          I : Integer;
11
       end record;
12
13
       function Init return Rec is
14
         ((I => 0));
15
16
   end Unknown Discriminants;
17
```

Listing 88: show_constructor_function.adb

```
1 with Unknown_Discriminants;
2 use Unknown_Discriminants;
3
4 procedure Show_Constructor_Function is
5 R : Rec := Init;
6 begin
7 null;
8 end Show_Constructor_Function;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Discriminants.

⇔Object_Declaration

MD5: 1cee0c4b883b3a0c25fae0a5111db2a8
```

In the Show_Constructor_Function procedure from this example, we call the Init function to initialize the R object in its declaration (of Rec type). Note that for this specific type, this is the only possible way to declare the R object. In fact, compilation fails if we write R : Rec;.

Using a private type with unknown discriminants is an important Ada idiom, as we gain extra control over its initialization. For example, if we have to ensure that certain components of the private record are initialized when an object is being declared, we can perform this initialization in the Init function — instead of just hoping that an initialization function is called for this object at some point. Also, if further information is needed to initialize an object, we can add parameters to the Init function, thereby forcing the user to provide this information.

For even more control over objects, we can use *limited types with unknown discriminants* (page 812).

4.6.2 Partial and full view

As we've just seen, if we declare a type with an unknown discriminant part, we can only use it in the partial view. In the full view. we cannot use an unknown discriminant part, but have to use either no discriminants or known discriminants. For example:

Listing 89: unknown_discriminants.ads

```
1 package Unknown_Discriminants is
2
3 type Rec_No_Discr (<>) is private;
4
5 type Rec_Known_Discr (<>) is private;
```

```
<sup>6</sup>
7
private
<sup>8</sup>
9
type Rec_No_Discr is null record;
10
11
type Rec_Known_Discr
12
(L : Positive) is null record;
13
```

```
14 end Unknown_Discriminants;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Partial_

⊸Full_View

MD5: 3d37dcc9d1b12bf9a189cf515b168430
```

In this example, Rec_No_Discr has no discriminants in its full view, while Rec_Known_Discr has the discriminant L.

In addition, the full view can be an (unconstrained) array type as well:

Listing 90: unknown_discriminants.ads

```
package Unknown_Discriminants is
1
2
      type Arr (<>) is private;
3
4
   private
5
6
      type Arr is
7
         array (Positive range <>)
8
           of Integer;
9
10
   end Unknown_Discriminants;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Partial_

GFull_View

MD5: dle0f60048c6ca6bcf863a8c0cf68314
```

Here, the full view of Arr is an array type.

In the Ada Reference Manual

• 3.7 Discriminants¹⁰¹

4.6.3 Derived types

As expected, we can derive from types with unknown discriminants. Consider the following package:

Listing 91: unknown_discriminants.ads

```
package Unknown_Discriminants is
```

```
type Rec (<>) is private;
```

3

```
<sup>101</sup> http://www.ada-auth.org/standards/12rm/html/RM-3-7.html
```

```
4
5 private
6
7 type Rec is null record;
8
9 end Unknown_Discriminants;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_Type
MD5: 948e7c7ecd00915fa23a98cbaf2bbcbe

We can then declare the Derived_Rec type:

Listing 92: unknown discriminants-children.ads

```
1 package Unknown_Discriminants.Children is
2
3 type Derived_Rec is
4 new Rec;
5
6 end Unknown_Discriminants.Children;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_Type
MD5: 1fa7e905c794d48bf6c76ff51e1abd8d
```

Note that Derived_Rec has unknown discriminants, even though we're not explicitly using an unknown discriminant part ((<>)) in its declaration. (In fact, we're not allowed to use an unknown discriminant part in this case.) Therefore, declaring objects of this type directly isn't possible, just like the parent type Rec:

Listing 93: show object declaration.adb

```
1 with Unknown_Discriminants.Children;
2 use Unknown_Discriminants.Children;
3
4 procedure Show_Object_Declaration is
5 A : Derived_Rec;
6 begin
7 null;
8 end Show_Object_Declaration;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_Type
MD5: 5d0b8980e6f60595b9de8a2ea8fa2132

Build output

```
show_object_declaration.adb:5:08: error: unconstrained subtype not allowed (need_
__initialization)
gprbuild: *** compilation phase failed
```

Deriving from tagged types

We can also derive from tagged types with unknown discriminants. Consider the following package:

Listing 94: unknown_discriminants.ads

```
1 package Unknown_Discriminants is
2
3 type Rec (<>) is tagged private;
4
5 private
6
7 type Rec is tagged null record;
8
9 end Unknown_Discriminants;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_

Gagged_Type

MD5: ef66d098df1c93495bf5f6c6ac86f203
```

We can derive from the Rec type. In this case, however, we can use an unknown discriminant part, a known discriminant part, or no discriminants:

Listing 95: unknown_discriminants-children.ads

```
package Unknown Discriminants.Children is
1
2
      type Derived Rec Unknown Discr (<>) is
3
         new Rec with private;
4
5
      type Derived_Rec_Known_Discr (L : Positive) is
6
         new Rec with private;
7
8
      type Derived_Rec_No_Discr is
9
         new Rec with private;
10
11
   private
12
13
      type Derived Rec Unknown Discr is
14
         new Rec with null record;
15
16
      type Derived_Rec_Known_Discr (L : Positive) is
17
         new Rec with null record;
18
19
      type Derived_Rec_No_Discr is
20
         new Rec with null record;
21
22
   end Unknown Discriminants.Children;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_

Gagged_Type

MD5: 98583f0b39c6f8bea49d1781844bb33e
```

In this example, we declare Derived_Rec_Unknown_Discr with an unknown discriminant part, Derived_Rec_Known_Discr with a known discriminant part, and Derived_Rec_No_Discr with no discriminants.

As expected, Derived_Rec_Unknown_Discr has unknown discriminants because it has an unknown discriminant part. In the case of Derived_Rec_No_Discr, which has no discriminants, we're deriving the unknown discriminants of Rec, so it also has unknown discriminants. In contrast, because Derived_Rec_Known_Discr has a known discriminant part, those discriminants are overriding the unknown discriminants of the parent type Rec. Therefore, we can declare objects of Derived_Rec_Known_Discr type without explicit initialization:

Listing 96:	show	_object_	decla	aration	.adb
-------------	------	----------	-------	---------	------

```
with Unknown Discriminants.Children;
1
   use Unknown_Discriminants.Children;
2
3
   procedure Show Object Declaration is
4
      A : Derived_Rec_Unknown_Discr;
5
      -- ERROR: unknown discriminants
6
                  because of the type's
       - -
7
                 unknown discriminant part
       - -
8
9
      B : Derived_Rec_Known_Discr (1);
10
      -- OK: known discriminants
11
12
      C : Derived Rec No Discr;
13
          ERROR: unknown discriminants
      - -
14
       - -
                  because of parent type's
15
       - -
                  unknown discriminant part
16
   begin
17
      null;
18
   end Show_Object_Declaration;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_
```

MD5: 91f6ae2abf88976833d0e4eff02d4c40

Build output

As we can see, we can only directly declare objects of type Derived_Rec_Known_Discr because it has known discriminants, while the other two derived types have unknown discriminants — which are explicitly specified (Derived_Rec_Unknown_Discr) or implicitly derived from the parent (Derived_Rec_No_Discr).

Note that the parent type Rec had a requirement for explicit initialization. By using known discriminants in the declaration of Derived_Rec_Known_Discr, we're removing this requirement for the derived type.

The contrary is also true: we can derive a type with known discriminants and use an unknown discriminant part:

Listing 97: unknown_discriminants-children-grand.ads

```
package Unknown Discriminants.Children.Grand is
1
2
      type Grand Rec Unknown Discr (<>) is
3
        new Derived Rec Known Discr (1)
4
          with private;
5
6
   private
7
8
      type Grand Rec Unknown Discr is
9
         new Derived Rec Known Discr (1)
10
```

with null record;

11 12

13 end Unknown_Discriminants.Children.Grand;

Listing 98: show_object_declaration.adb

```
with Unknown Discriminants.Children.Grand;
1
   use Unknown_Discriminants.Children.Grand;
2
3
   procedure Show Object Declaration is
4
      A : Grand_Rec_Unknown_Discr;
5
      -- ERROR: unknown discriminants
6
                 because of the type's
      - -
7
      - -
                  unknown discriminant part
8
   begin
9
      null:
10
   end Show Object Declaration;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_

⊸Tagged_Type

MD5: 0e931df8294cee1f49b187c43614aa20
```

Build output

In this example, Grand_Rec_Unknown_Discr has unknown discriminants and requires explicit initialization, even though its parent type Derived_Rec_Known_Discr has known discriminants.

1 In the Ada Reference Manual

• 3.7 Discriminants¹⁰²

4.7 Unconstrained subtypes

A subtype is called an unconstrained subtype if its type has unknown discriminants. Consider a simple Rec type:

Listing 9	99:	unknown	discriminants.ads

```
1 package Unknown_Discriminants is
2
3 type Rec (<>) is private;
4
5 private
6
```

```
<sup>102</sup> http://www.ada-auth.org/standards/12rm/html/RM-3-7.html
```

7 type Rec is null record;

```
8
9
```

end Unknown_Discriminants;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.

⊲Unconstrained_Subtype

MD5: 948e7c7ecd00915fa23a98cbaf2bbcbe
```

A subtype of Rec type is unconstrained:

Listing 100: unknown_discriminants-children.ads

```
package Unknown_Discriminants.Children is
subtype Rec_Unconstrained is Rec;
end Unknown_Discriminants.Children;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.

GUnconstrained_Subtype

MD5: 6b76b6a94d8c9487dbeea3256d5de01f
```

In this example, Rec_Unconstrained is an unconstrained subtype because it's derived from the Rec type. We can verify this by triggering a compilation error:

Listing 101: show_object_declaration.adb

```
with Unknown_Discriminants.Children;
use Unknown_Discriminants.Children;
procedure Show_Object_Declaration is
A : Rec_Unconstrained;
begin
null;
end Show_Object_Declaration;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Derived_Type
MD5: 442fab4d174de31f27d0de56bf9b8422

Build output

```
show_object_declaration.adb:5:08: error: "Rec_Unconstrained" is undefined
gprbuild: *** compilation phase failed
```

In addition, if we declare a subtype based on a type that allows range, index, or discriminant constraints, but we don't constraint the subtype, this subtype is also considered an unconstrained subtype. For example:

Listing 102: unconstrained_subtypes.ads

```
1 package Unconstrained_Subtypes is
2
3 type Arr is
4 array (Positive range <>) of
5 Integer;
```

```
6
       type Rec (L : Positive) is
7
         null record;
8
9
       subtype Arr_Sub is Arr;
10
11
       -- no constraints
12
13
       subtype Rec_Sub is Rec;
14
15
       - -
          no constraints
16
17
   end Unconstrained_Subtypes;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Unknown_Discriminants.Other_

GUnconstrained_Subtypes

MD5: 3ebc2eb371472dc76eb543b4633e59b3
```

In this example, Arr_Sub and Rec_Sub are unconstrained subtypes.

In the Ada Reference Manual

3.2 Types and Subtypes¹⁰³

4.8 Variant parts

We've introduced variant records back in the Introduction to Ada course¹⁰⁴. In simple terms, a variant record is a record with discriminants that allows for varying its structure. Basically, it's a record containing a **case** statement that specifies which record components exist for each discriminant value. For example:

```
package Devices is
1
2
       type Device State is
3
         (Off, On);
4
5
       type Device Info is
6
       record
7
          V : Float;
8
       end record;
9
10
       type Device (State : Device State := Off) is
11
       record
12
          case State is
13
             when Off =>
14
                 null;
15
             when On =>
16
                 Info : Device_Info;
17
          end case;
18
       end record;
19
```

¹⁰³ http://www.ada-auth.org/standards/12rm/html/RM-3-2.html
¹⁰⁴ https://learn.adacore.com/courses/intro-to-ada/chapters/more_about_records.html#
intro-ada-variant-records

20

21 end Devices;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Simple_Device
MD5: 3b63a63aef1d9cb00be870c831829158
```

The Device type from this example has a variant part. Depending on the value of the State discriminant, it can be either a null record (when State is Off) or have the Info component (when State is On).

Let's look at a test application for the Devices package:

Listing 104: show_device.adb

```
with Devices; use Devices;
1
2
   procedure Show_Device is
3
      D : Device;
4
      D_Off : Device (Off);
5
      D_On : Device (On);
6
7
   begin
      D := D_0ff;
8
      -- OK!
9
10
      D := D On;
11
      -- OK!
12
13
      D_Off := D_On;
14
15
       -- CONSTRAINT_ERROR!
16
17
      D_On := D_Off;
18
19
       -- CONSTRAINT ERROR!
20
   end Show_Device;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Simple_Device
MD5: alle2739131f435e8428a5e2a9a478e7
```

Build output

Runtime output

raised CONSTRAINT_ERROR : show_device.adb:14 discriminant check failed

As we've discussed *previously* (page 193), when we set the values for the discriminants of a type in the object declaration, we're constraining the objects. If the discriminants of two objects don't match, the Constraint_Error exception is raised at runtime because the *discriminant check* (page 515) fails. Therefore, in the Show_Device procedure, because D_Off and D_On are constrained and have different values for the State discriminant, we cannot assign them to each other. In contrast, because D wasn't constrained at its declaration, we can assign objects with different discriminants (such as D_Off and D_On) to it.

Note that the variant part of a record can be more complex. For example, we could have an additional discriminant and use it in the variant part:

Listing 10)5: de	vices.ads
------------	--------	-----------

```
package Devices is
1
2
       type Device State is
3
         (Off, On);
4
5
       type Device_Info is
6
       record
7
          V : Float;
8
       end record;
9
10
       type Device (State : Device_State;
11
                      Boost : Boolean) is
12
       record
13
          case State is
14
             when Off =>
15
                 null;
16
             when On =>
17
                 Info : Device_Info;
18
                 case Boost is
19
                    when False =>
20
                        null;
21
                     when True =>
22
                        Factor : Float;
23
                 end case;
24
          end case;
25
       end record;
26
27
   end Devices;
28
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Device_Boost
MD5: 4c5e84ccebca9e4ef5e2d6d131ba0e6a

In this version of the Devices package, we introduced a *boost button* as a discriminant (Boost) and an associated boost factor component (Factor) in the variant part.

In the remaining parts of this section, we discuss a couple of details about variant records.

In the Ada Reference Manual

• 3.8.1 Variant Parts and Discrete Choices¹⁰⁵

¹⁰⁵ http://www.ada-auth.org/standards/12rm/html/RM-3-8-1.html

4.8.1 Discriminant type and value coverage

The subtype of discriminants used in the variant part must be of a discrete type — it cannot be of an access or a floating-point type, for example. Also, all possible values of the subtype of each discriminant must be covered in the case statement of the variant part. For example, consider the following variant record:

Listing 106: subtype_coverage.ads

```
package Subtype_Coverage is
1
2
      type Var_Rec (Value : Integer) is
3
       record
4
          case Value is
5
             when 0 .. 100 =>
6
                I : Integer;
7
8
             -- ERROR: missing values!
9
          end case;
10
      end record;
11
12
   end Subtype_Coverage;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Coverage
MD5: 084a468bc8d6f63d21f804c9ddc70622

Build output

subtype_coverage.ads:5:07: error: missing case values: -16#8000_0000# .. -1
subtype_coverage.ads:5:07: error: missing case values: 101 .. 16#7FFF_FFFF#
gprbuild: *** compilation phase failed

This package cannot be compiled because, in the variant part, we're only covering values for the Value discriminant in the range between 0 and 100. To fix this compilation error, we have to cover all values instead. For example:

```
package Subtype_Coverage is
1
2
      type Var Rec (Value : Integer) is
3
      record
4
          case Value is
5
             when Integer'First .. -1 =>
6
                null;
7
             when 0 .. 100 =>
8
                I : Integer;
9
             when 101 .. Integer'Last =>
10
                null;
11
          end case;
12
      end record;
13
14
  end Subtype_Coverage;
15
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Coverage
MD5: 9dfa0dfc3d3e11394a79b1ab6b61bafc

Of course, specifying all possible values can be difficult. As an alternative, we could simplify the case statement by just using **others** as a discrete choice that encompasses all values

that haven't been specified earlier in the case statement:

Listing 108: subtype coverage.ads

```
package Subtype Coverage is
1
2
      type Var Rec (Value : Integer) is
3
       record
4
          case Value is
5
             when 0 .. 100 =>
6
                I : Integer;
7
             when others =>
8
                 null;
9
          end case;
10
      end record;
11
12
   end Subtype_Coverage;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Coverage
MD5: 0b28038d5137de702cb5b8e875fadefa
```

By using **when others** => ... in this last example, we ensure that all values have been covered.

4.8.2 Record size

When declaring an object, the values we select for the discriminants related to the variant part have an impact on the overall size of that object — in fact, it may be smaller or bigger depending on this selection. Let's see an example:

```
Listing 109: variant_records.ads
```

```
package Variant_Records is
1
2
      type Simple_Record
3
        (Extended : Boolean := False) is
4
      record
5
          V : Integer;
6
          case Extended is
7
             when False =>
8
                null;
9
             when True =>
10
                V_Float : Float;
11
          end case;
12
      end record;
13
14
   end Variant Records;
15
```

Listing 110: show variant rec size.adb

```
with Ada.Text I0;
                          use Ada.Text I0;
1
2
  with Variant_Records; use Variant_Records;
3
4
   procedure Show_Variant_Rec_Size is
5
      SR_No_Ext : Simple_Record
6
                     (Extended => False);
7
                : Simple Record
      SR Ext
8
                     (Extended => True);
9
```

```
SR
                  : Simple_Record;
10
   begin
11
      Put Line ("SR No Ext'Size : "
12
                 & SR_No_Ext'Size'Image
13
                 & " bits");
14
       Put_Line ("SR_Ext'Size :
15
                 & SR_Ext'Size'Image
16
                 & " bits");
17
      Put_Line ("SR'Size :
18
                 & SR'Size'Image
19
                 & " bits");
20
   end Show_Variant_Rec_Size;
21
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Size
MD5: 4aaf10924e7469d000cefeb70a69a2fa

Build output

Runtime output

SR_No_Ext'Size : 64 bits
SR_Ext'Size : 96 bits
SR'Size : 96 bits

As we can confirm when we run this application, the choice for the discriminant has an impact on the size of the object. In the case of the SR_No_Ext object, setting the Extended discriminant to **False** excludes the V_Float component. For the SR_Ext object, on the other hand, we include the V_Float component. Therefore, on a typical PC, the size of SR_No_Ext is 8 bytes (4 bytes for the Extended discriminant and 4 bytes for the V component), while the size of SR_Ext is 12 bytes (i.e., additional 4 bytes for the V_Float component).

In the case of SR, because the object isn't constrained, the size of the object is 12 bytes on a typical PC — the same size as SR_Ext. This is because SR has to account for the case when all components must be available, even though the Extended discriminant is set to **False** by default. Remember that an assignment such as SR := SR_Ext is valid, so enough memory must be available to ensure that the assignment is performed correctly.

This principle applies to more complicated variant records. For example:

```
Listing 111: variant_records.ads
```

```
package Variant Records is
1
2
      type Simple Record
3
         (Extended : Boolean := False;
4
          Extended 2 : Boolean := False) is
5
      record
6
         V : Integer;
7
         case Extended is
8
             when False =>
9
                case Extended 2 is
10
                   when False =>
11
```

```
null;
12
                     when True =>
13
                        V_Int_2 : Integer;
14
                        V_Int_3 : Integer;
15
                 end case;
16
              when True =>
17
                 V_Float : Float;
18
                 case Extended 2 is
19
                    when False =>
20
                        null;
21
                     when True =>
22
                        V_Float_2 : Float;
23
                 end case;
24
25
          end case;
26
       end record;
27
   end Variant_Records;
28
```

Listing 112: show_variant_rec_size.adb

```
with Ada.Text_I0;
                           use Ada.Text_I0;
1
2
   with Variant_Records; use Variant_Records;
3
4
   procedure Show Variant Rec Size is
5
      SR : Simple_Record;
6
   begin
7
      Put Line ("SR'Size : "
8
                 & SR'Size'Image
9
                 & " bits");
10
   end Show_Variant_Rec_Size;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Size
MD5: 5f0ac936a5fee50cbe88a7b863a1a550
```

Build output

Runtime output

SR'Size : 128 bits

In this example, the size of SR is 16 bytes on a typical PC. This accounts for 4 bytes for the discriminants Extended and Extended_2, and 4 bytes for each of the 3 components that are being taken into account for the worst case:

- components V, V_Int_2 and V_Int_3 when we set Extended => False, Extended_2 => True;
- components V, V_Float and V_Float_2 when we set Extended => True, Extended_2
 => True.

Note that a memory block is shared between the V_Int_2 and V_Int_3 components from the first worst case, and V_Float and the V_Float_2 components from the second worst case. As we can see, the compiler will typically optimize the size of a record as much as possible by assessing which components are really needed for the worst case.

Also, as we discussed previously, we can use *unchecked unions* (page 117) in combination with variant records, which has an impact on the object size.

4.8.3 Ensuring valid information

We can use variant parts to prevent invalid information from being used. Let's look again at the Device type from the previous code example:

```
type Device (State : Device_State) is
record
   case State is
     when Off =>
        null;
     when On =>
        Info : Device_Info;
   end case;
end record;
```

For the sake of this example, we could say that a device that is turned off doesn't have any valuable information. Therefore, the device information stored in the Info component of the Device type is only valid if the device is turned on. Thus, if the device is turned off (i.e., Device_State = Off), we should prevent the application from processing device information that is probably incorrect. Let's extend the previous code example to accommodate this requirement:

Listing 113: devices.ads

```
package Devices is
1
2
       type Device_State is
3
         (Off, On);
4
5
       type Device
6
         (State : Device_State := Off) is
7
           private;
8
9
       procedure Turn_Off (D : in out Device);
10
11
       procedure Turn_On (D : in out Device);
12
13
       type Device Info is
14
       record
15
          V : Float;
16
       end record;
17
18
       function Current_Info (D : Device)
19
                                return Device_Info;
20
21
   private
22
23
       type Device (State : Device_State := Off) is
24
       record
25
          case State is
26
             when Off =>
27
                null;
28
             when 0n =>
29
                 Info : Device Info;
30
          end case;
31
       end record;
32
33
       Device_Off : constant Device :=
34
```

```
35 (State => Off);
36 37 Device_On : constant Device :=
38 (State => On,
39 others => <>);
40 41 end Devices;
```

Listing 114: devices.adb

```
package body Devices is
1
2
      procedure Turn_Off (D : in out Device) is
3
      beain
4
          D := Device Off;
5
      end Turn_Off;
6
7
      procedure Turn On (D : in out Device) is
8
      begin
9
          D := Device On;
10
      end Turn_On;
11
12
      function Current_Info (D : Device)
13
                                return Device_Info is
14
         (D.Info);
15
16
   end Devices;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Device
MD5: e03db406de3550865dd99986d2c71145
```

Let's create a test application called Show_Device that makes use of this device by turning it on and off, and by retrieving information from it:

Listing 115: show_device.adb

```
with Devices; use Devices;
1
2
   procedure Show_Device is
3
      D : Device;
4
      I : Device_Info;
5
   begin
6
      Turn On (D);
7
      I := Current_Info (D);
8
9
      Turn_Off (D);
10
11
       -- The following call raises
12
       -- an exception at runtime
13
          because D is turned off.
14
       - -
      I := Current_Info (D);
15
   end Show Device;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Device
MD5: cba0100ad5bbb2b6bf00d0847a700271
```

Runtime output

raised CONSTRAINT_ERROR : devices.adb:15 discriminant check failed

In this example, by using the variant part, we're preventing information retrieved by an inappropriate call to the Current_Info function from being used elsewhere in the application. In fact, if the device is turned off, a call to Current_Info raises the Constraint_Error exception because the Info component isn't accessible. We see that effect in the Show_Device procedure: the call to Current_Info fails (by raising an exception) when the device has just been turned off.

To avoid exceptions at runtime, we must check the device's state before calling Current_Info:

with Devices; use Devices;
<pre>procedure Show_Device is</pre>
D : Device;
I : Device_Info;
begin
Turn_On (D);
<pre>if D.State = On then I := Current_Info (D); end if;</pre>
<pre>Turn_Off (D);</pre>
<pre>if D.State = On then I := Current_Info (D); end if; end Show_Device;</pre>

Listing 116: show device.adb

Code block metadata

1 W

3 4 5

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Device
MD5: 62230848af720b156f22c96d59f772d2
```

Now, no exception is raised, as we only retrieve information from the device when it is turned on — that is, we only call the Current_Info function when the State discriminant of the object is set to On.

4.8.4 Extending record types

We can use variant parts as a means to extend record types. This can be viewed as a static approach to implement type extension — similar to type extension via tagged types, but with clear differences.

Let's say we have a sensor, and we implement a package called Sensors that interfaces with that sensor:

Listing 117: sensors.ads

```
package Sensors is
type Sensor is private;
type Sensor_Info is
record
Info_1 : Float := 0.0;
```

```
end record;
8
9
       function Current_Info (S : Sensor)
10
                                 return Sensor_Info;
11
12
       procedure Display (SI : Sensor_Info);
13
14
   private
15
16
       type Sensor is null record;
17
18
   end Sensors;
19
```

Listing 118: sensors.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Sensors is
3
4
       function Current Info (S : Sensor)
5
                                return Sensor_Info is
6
         ((Info_1 => 4.0));
7
8
       - -
       -- NOTE: we're returning dummy
9
                 information!
10
11
      procedure Display (SI : Sensor_Info) is
12
      begin
13
          Put_Line ("Info_1 : "
14
                    & SI.Info_1'Image);
15
      end Display;
16
17
   end Sensors;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Sensors
MD5: 140a0d9cbca023de875417409c3f67d9

The Sensor type from the Sensors package has two subprograms: the Current_Info function and the Display procedure. We use those subprograms in the Show_Sensors procedure below:

Listing 119: show_sensors.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
  with Sensors;
                     use Sensors;
3
4
  procedure Show_Sensors is
5
      S1 : Sensor;
6
  begin
7
      Display (Current_Info (S1));
8
  end Show_Sensors;
9
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Sensors
MD5: 93aa76da463fea9b4483ed97fa8bcf64
```

Runtime output

Info_1 : 4.00000E+00

Now, let's assume that a new model of this sensor is available, and it has additional features — e.g., it provides additional information to the user. If we wanted to update the application to be able to handle this new model of the sensor without removing support for the original model, we could convert the Sensor_Info type to a tagged type and derive a Sensor_Info_V2 type from it. (We would probably have to implement a Sensor_V2 type derived from the Sensor type as well.)

Alternatively, we could add a variant part to the Sensor_Info type to store the additional information. For example:

Listing 120: sensors.ads

```
package Sensors is
1
2
       type Sensor_Model is (Sensor_V1,
3
                               Sensor V2);
4
5
       type Sensor
6
         (Model : Sensor Model := Sensor V1) is
7
           private;
8
9
       type Sensor_Info
10
         (Model : Sensor_Model := Sensor_V1) is
11
       record
12
          Info_1 : Float := 0.0;
13
          case Model is
14
             when Sensor_V1 =>
15
                null;
16
             when Sensor V2 =>
17
                 Info_2 : Float := 0.0;
18
          end case;
19
20
       end record;
21
       function Current_Info (S : Sensor)
22
                                 return Sensor_Info;
23
24
       procedure Display (SI : Sensor_Info);
25
26
   private
27
28
       type Sensor
29
         (Model : Sensor Model := Sensor V1) is
30
31
           null record;
32
   end Sensors;
33
```

Listing 121: sensors.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Sensors is
3
4
      function Current_Info (S : Sensor)
5
                               return Sensor_Info is
6
      begin
7
         - -
             Using dummy info for the information
8
9
          -- returned by the function
10
         case S.Model is
             when Sensor_V1 =>
11
```

```
return ((Model => Sensor_V1,
12
                           Info_1 => 4.0));
13
             when Sensor V2 =>
14
                 return ((Model => Sensor_V2,
15
                           Info_1 => 8.0,
16
                           Info_2 => 6.0));
17
          end case:
18
       end Current_Info;
19
20
       procedure Display (SI : Sensor_Info) is
21
       begin
22
                             1.1
          Put_Line ("Model
23
                     & SI.Model'Image);
24
          Put_Line ("Info_1 : "
25
                     & SI.Info_1'Image);
26
          if SI.Model = Sensor_V2 then
27
             Put_Line ("Info_2 : "
28
                        & SI.Info_2'Image);
29
          end if;
30
       end Display;
31
32
   end Sensors;
33
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Sensors
MD5: 74198e928e3dfa3a7a7f2786971da8a7

In this new version of the Sensors package, the Model discriminant was added to the Sensor_Info type. If the model is set to version 2 for a specific sensor (i.e., Model = Sensor_V2), a new component (Info_2) is available.

The Current_Info and Display subprograms have been adapted to take this new model into account. In the Current_Info function, we return information for the newer model of the sensor. In the Display procedure, we display the additional information provided by the newer model.

Note that the original test application that makes use of the sensor (Show_Sensors) doesn't require any adaptation:

Listing 122: show sensors.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
  with Sensors;
                      use Sensors;
3
4
  procedure Show_Sensors is
5
      S1 : Sensor;
6
  begin
7
      Display (Current_Info (S1));
8
  end Show_Sensors;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Sensors
MD5: 93aa76da463fea9b4483ed97fa8bcf64

Runtime output

Info_1 : 4.00000E+00

Because we have a default value for the discriminant of the Sensor type, we're essentially

making the type *backwards-compatible*, so that users of this type don't have to adapt their code after the update to the Sensors package. Of course, we don't have *binary backwards-compatibility* because the size of the type (Sensor_Info'Size) increases.

Of course, in our test application, we can also use the new model of that sensor:

```
Listing 123: show_sensors.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
   with Sensors:
                    use Sensors;
3
4
   procedure Show Sensors is
5
      S1 : Sensor;
6
      S2 : Sensor (Sensor V2);
7
   begin
8
      Display (Current Info (S1));
9
      Display (Current Info (S2));
10
  end Show Sensors;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Records.Variant_Parts.Sensors
MD5: 347d272cddbacf7bf2987aa23014ff0b

Runtime output

Model : SENSOR_V1 Info_1 : 4.00000E+00 Model : SENSOR_V2 Info_1 : 8.00000E+00 Info_2 : 6.00000E+00

In the updated version of the Show_Sensors procedure, we're now using both old and new versions of the sensor.

4.9 Per-Object Expressions

In record type declarations, we might want to define a component that makes use of a *name* (page 5) that refers to a *discriminant* (page 188) of the record type, or to the record type itself. An expression where we use such a name is called a per-object expression.

The term "per-object" comes from the fact that, in the component definition, we're referring to a piece of information that will be known just when creating an object of that type. For example, if the per-object expression refers to a discriminant of a type T, the actual value of that discriminant will only be specified when we declare an object of type T. Therefore, the component definition is specific for that individual object — but not necessarily for other objects of the same type, as we might use different values for the discriminant.

The constraint that contains a per-object expression is called a per-object constraint. The actual constraint of that component isn't completely known when we declare the record type, but only later on when an object of that type is created. (Note that the syntax of a constraint includes the parentheses or the keyword **range**.)

In addition to referring to discriminants, per-object expressions can also refer to the record type itself, as we'll see later.

Let's start with a simple record declaration:

Listing 124: rec_per_object_expressions.ads

```
package Rec Per Object Expressions is
1
2
       type Stack (S : Positive) is private;
3
4
   private
5
6
       type Integer Array is
7
         array (Positive range <>) of Integer;
8
9
       type Stack (S : Positive) is record
10
          Arr : Integer_Array (1 .. S);
11
                                   ~~~~~
12
          - -
          - -
13
          - -
                                        S
14
           - -
                                        ^
15
                 Per-object expression
          - -
16
          - -
17
                                 (1 .. S)
          - -
18
                                   ~~~~~
          - -
19
          - -
                  Per-object constraint
20
21
          Top : Natural := 0;
22
       end record;
23
24
   end Rec_Per_Object_Expressions;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Per_Object_Expressions.Per_Object_

⇔Expression

MD5: e4012454ea886fd429d82159b8d344b7
```

In this example, we see the Stack record type with a discriminant S. In the declaration of the Arr component of the that type, S is a per-object expression, as it refers to the S discriminant. Also, $(1 \dots S)$ is a per-object constraint.

Let's look at another example using anonymous access types (page 711):

Listing 125: rec_per_object_expressions.ads

```
package Rec_Per_Object_Expressions is
1
2
      type T is private;
3
4
       type T_Processor (Selected_T : access T) is
5
         private;
6
7
   private
8
9
      type T is null record;
10
11
      type T Container (Selected T : access T) is
12
         null record;
13
14
      type T Processor (Selected T : access T) is
15
       record
16
          E : T Container (Selected T);
17
          - -
18
          - -
                             Selected T
19
                             ~~~~~~~
20
          - -
```

```
- -
                  Per-object expression
21
           - -
22
           - -
                              (Selected T)
23
           - -
24
                   Per-object constraint
           - -
25
       end record;
26
27
   end Rec_Per_Object_Expressions;
28
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Per_Object_Expressions.Per_Object_

→Expression_Access_Discriminant

MD5: 8b404688be1e103773c28a6977785836
```

Let's focus on the T_Processor type from this example. The Selected_T discriminant is being used in the definition of the E component. The per-object constraint is (Selected_T).

Finally, per-object expressions can also refer to the record type we're declaring. For example:

Listing 126: rec_per_object_expressions.ads

```
package Rec_Per_Object_Expressions is
1
2
      type T is limited private;
3
4
   private
5
6
       type T_Processor (Selected_T : access T) is
7
         null record;
8
9
       type T is limited record
10
          E : T_Processor (T'Access);
11
          - -
12
          - -
                             T'Access
13
                             ~~~~~
          - -
14
          -- Per-object expression
15
          - -
16
          - -
                            (T'Access)
17
                             ~~~~~
          - -
18
          - -
              Per-object constraint
19
      end record;
20
21
  end Rec_Per_Object_Expressions;
22
```

Code block metadata

In this example, when we write T'Access within the declaration of the T record type, the actual value for the Access attribute will be known when an object of T type is created. In that sense, T'Access is a per-object expression — (T'Access) is the corresponding per-object constraint.

Note that T'Access is referring to the type within a type definition. This is generally treated as a reference to the object being created, the so-called *current instance*.

1 In the Ada Reference Manual

3.8 Record Types¹⁰⁶

4.9.1 Default value

We can also use per-object expressions to calculate the default value of a record component:

Listing 127: rec per object expressions.ads

```
package Rec Per Object Expressions is
1
2
       type T (D : Positive) is private;
3
4
   private
5
6
       type T (D : Positive) is record
7
          V : Natural := D - 1;
8
9
          - -
                Per-object expression
10
          - -
11
          S : Natural := D'Size;
12
                             ~~~
13
          - -
                 Per-object expression
14
15
       end record;
16
   end Rec_Per_Object_Expressions;
17
```

Code block metadata

Here, we calculate the default value of V using the per-object expression D -1, and the default of value of S using the per-object D'Size.

The default expression for a component of a discriminated record can be an arbitrary perobject expression. (This contrasts with *important restrictions* (page 244) that exist for perobject constraints, as we discuss later on.) Such expressions might include function calls or uses of any defined operator. For this reason, the following code example is accepted by the compiler:

Listing 128: rec_per_object_expressions.ads

```
package Rec_Per_Object_Expressions is
1
2
      type Stack (S : Positive) is private;
3
4
   private
5
6
      type Integer Array is
7
         array (Positive range <>) of Integer;
8
9
      type Stack (S : Positive) is record
10
         Arr : Integer Array (1 .. S);
11
12
          Top : Natural := 0;
13
```

(continues on next page)

¹⁰⁶ http://www.ada-auth.org/standards/22rm/html/RM-3-8.html

```
14
          Overflow_Warning : Positive
15
            := S * 9 / 10;
16
                ~~~~~~~
          - -
17
              Per-object expression
          - -
18
              using computation for
19
          - -
              the default expression.
          - -
20
       end record
21
        with
22
          Dynamic Predicate =>
23
             Overflow_Warning in
24
                (S + 1) / 2 .. S - 1;
25
26
          - -
               (S + 1) / 2
27
          - -
                ^^^^^
28
          - -
               Per-object expression
          - -
29
          - -
               using computation.
30
          - -
31
                                S - 1
          - -
32
                                ~~~~
          - -
33
              Per-object expression
          - -
34
          - -
              using computation.
35
36
   end Rec_Per_Object_Expressions;
37
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Per_Object_Expressions.Per_Object_

⇔Expression_Computation

MD5: 6783568fd3e76a85ca7c1cc65ba023c5
```

In this example, we can identify multiple per-object expressions that use a computation: S $\ast~9$ / 10, (S + 1) / 2, and S - 1.

4.9.2 Restrictions

There are some important restrictions on per-object constraints:

- 1. Per-object range constraints such as 1 . . T'Size are not allowed.
 - For example, the following code example doesn't compile:

Listing 129: rec_per_object_expressions.ads

```
package Rec_Per_Object_Expressions is
1
2
       type Bit Field is
3
         array (Positive range <>) of Boolean
4
           with Pack;
5
6
      type T is record
7
          Arr : Bit_Field (1 .. T'Size);
8
9
          - -
          -- ERROR: per-object range constraint
10
                     using the Size attribute
          - -
11
                     is illegal.
12
          - -
       end record;
13
14
  end Rec Per Object Expressions;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Per_Object_

←Expressions.Per_Object_Expression_Range_Constraint

MD5: c2ac9588c1d1adac8c584a0e36a81342
```

Build output

- 2. Within a per-object index constraint or discriminant constraint, each per-object expression must be the name of a discriminant directly, without any further computation.
 - Therefore, we're allowed to write (1 .. S) as we've seen in a previous example
 —. However, writing (1 .. S 1) would be illegal.
 - For example, the following adaptation to the previous code example doesn't compile:

Listing 130: rec_per_object_expressions.ads

```
package Rec_Per_Object_Expressions is
1
2
       type Stack (S : Positive) is private;
3
4
   private
5
6
       type Integer Array is
7
         array (Natural range <>) of Integer;
8
9
       type Stack (S : Positive) is record
10
          Arr : Integer_Array (0 .. S - 1);
11
                                       ~ ~ ~ ~ ~
12
          -- ERROR: computation in per-object
13
                      expression is illegal.
          - -
14
15
          Top : Integer := -1;
16
       end record;
17
18
   end Rec_Per_Object_Expressions;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Records.Per_Object_

←Expressions.Per_Object_Expression_Range_Computation

MD5: 1224bb63f7953743d84a258226c35c50
```

Build output

In this example, using the computation S - 1 to specify the range of Arr isn't permitted. (Note that, *as we've seen before* (page 243), this restriction doesn't apply when the computation is used in a per-object expression that calculates the default value of a component.)

3. We can only use access attributes (T'Access and T'Unchecked_Access) in per-object constraints.

CHAPTER FIVE

AGGREGATES

5.1 Container Aggregates

\rm Note

This feature was introduced in Ada 2022.

A container aggregate is a list of elements — such as [1, 2, 3] — that we use to initialize or assign to a container. For example:

Listing 1: show_container_aggregate.adb

```
with Ada.Containers.Vectors;
1
2
   procedure Show Container Aggregate is
3
4
      package Float Vec is new
5
        Ada.Containers.Vectors (Positive, Float);
6
7
      V : constant Float Vec.Vector :=
8
             [1.0, 2.0, 3.0];
9
10
      pragma Unreferenced (V);
11
  begin
12
      null;
13
  end Show_Container_Aggregate;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Container_Aggregates.Simple_

GContainer_Aggregate

MD5: b54cd5800179d4016bbce5a9b10734f2
```

In this example, [1.0, 2.0, 3.0] is a container aggregate that we use to initialize a vector V.

We can specify container aggregates in three forms:

- as a null container aggregate, which indicates a container without any elements and is represented by the [] syntax;
- as a positional container aggregate, where the elements are simply listed in a sequence (such as [1, 2]);
- as a named container aggregate, where a key is indicated for each element of the list (such as [1 => 10, 2 => 15]).

Let's look at a complete example:

```
Listing 2: show_container_aggregate.adb
```

```
with Ada.Containers.Vectors;
1
2
   procedure Show Container Aggregate is
3
4
      package Float_Vec is new
5
        Ada.Containers.Vectors (Positive, Float);
6
7
       -- Null container aggregate
8
      Null V : constant Float Vec.Vector :=
9
                    [];
10
11
       -- Positional container aggregate
12
      Pos_V : constant Float_Vec.Vector :=
13
                    [1.0, 2.0, 3.0];
14
15
       -- Named container aggregate
16
      Named V : constant Float Vec.Vector :=
17
                    [1 \Rightarrow 1.0,
18
                     2 => 2.0,
19
                     3 => 3.0];
20
21
      pragma Unreferenced (Null_V, Pos_V, Named_V);
22
   begin
23
      null;
24
   end Show_Container_Aggregate;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Container_Aggregates.Simple_

GContainer_Aggregate

MD5: f00b21da1722669ae92bd5fe4a9a3966
```

In this example, we see the three forms of container aggregates. The difference between positional and named container aggregates is that:

for positional container aggregates, the vector index is implied by its position;

while

 for named container aggregates, the index (or key) of each element is explicitly indicated.

Also, the named container aggregate in this example (Named_V) is using an index as the name (i.e. it's an indexed aggregate). Another option is to use non-indexed aggregates, where we use actual keys — as we do in maps. For example:

Listing 3: show_named_container_aggregate.adb

```
with Ada.Containers.Vectors;
1
   with Ada.Containers.Indefinite_Hashed_Maps;
2
   with Ada.Strings.Hash;
3
   procedure Show_Named_Container_Aggregate is
5
6
      package Float Vec is new
7
        Ada.Containers.Vectors (Positive, Float);
8
9
      package Float Hashed Maps is new
10
        Ada.Containers.Indefinite Hashed Maps
11
          (Key_Type => String,
12
                         => Float,
           Element_Type
13
```

```
Hash
                              => Ada.Strings.Hash,
14
            Equivalent_Keys => "=");
15
16
       -- Named container aggregate
17
       -- using an index
18
       Indexed_Named_V : constant Float_Vec.Vector :=
19
                              [1 \implies 1.0,
20
                               2 => 2.0,
21
                               3 \implies 3.0];
22
23
       -- Named container aggregate
24
       - -
          using a key
25
       Keyed_Named_V : constant
26
         Float_Hashed_Maps.Map :=
27
           ["Key_1" => 1.0,
28
             "Key_2" => 2.0,
29
            "Key_3" => 3.0];
30
31
       pragma Unreferenced (Indexed Named V,
32
                              Keyed_Named_V);
33
   begin
34
       null;
35
   end Show_Named_Container_Aggregate;
36
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Container_Aggregates.Named_

→Container_Aggregate

MD5: 9d117543135e75e66801628ca29e32ef
```

In this example, Indexed_Named_V and Keyed_Named_V are both initialized with a named container aggregate. However:

 the container aggregate for Indexed_Named_V is an indexed aggregate, so we use an index for each element;

while

• the container aggregate for Keyed_Named_V has a key for each element.

Later on, we'll talk about the Aggregate aspect, which allows for defining custom container aggregates for any record type.

```
1 In the Ada Reference Manual
```

• 4.3.5 Container Aggregates¹⁰⁷

5.2 Record aggregates

We've already seen record aggregates in the Introduction to Ada¹⁰⁸ course, so this is just a brief overview on the topic.

As we already know, record aggregates can have positional and named component associations. For example, consider this package:

¹⁰⁷ http://www.ada-auth.org/standards/22rm/html/RM-4-3-5.html

¹⁰⁸ https://learn.adacore.com/courses/intro-to-ada/chapters/records.html#intro-ada-record-aggregates

```
Listing 4: points.ads
```

```
package Points is

type Point_3D is record
X, Y, Z : Integer;
end record;

procedure Display (P : Point_3D);
end Points;
```

Listing 5: points.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Points is
3
4
5
       procedure Display (P : Point_3D) is
6
       begin
          Put_Line ("(X => "
7
                     & Integer'Image (P.X)
8
          & ",");
Put_Line (" Y => "
9
10
                     & Integer'Image (P.Y)
11
                     & ",");
12
          Put_Line (" Z => "
13
                     & Integer'Image (P.Z)
14
                     & ")");
15
       end Display;
16
17
   end Points;
18
```

Code block metadata

We can use positional or named record aggregates when assigning to an object P of Point_3D type:

Listing 6: show_record_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show_Record_Aggregates is
3
      P : Point_3D;
4
   begin
5
      -- Positional component association
6
      P := (0, 1, 2);
7
8
      Display (P);
9
10
       -- Named component association
11
      P := (X => 3,
12
             Y => 4,
13
             Z => 5);
14
15
      Display (P);
16
   end Show_Record_Aggregates;
17
```

Code block metadata

Runtime output

(X => 0, Y => 1, Z => 2) (X => 3, Y => 4, Z => 5)

Also, we can have a mixture of both:

```
Listing 7: show_record_aggregates.adb
```

```
with Points; use Points;
1
2
   procedure Show_Record_Aggregates is
3
      P : Point_3D;
4
   begin
5
      -- Positional and named component associations
6
      P := (3, 4,
7
             Z => 5);
8
9
      Display (P);
10
   end Show_Record_Aggregates;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Pos_Named_

←Rec_Aggregates

MD5: 493a2a87b4b28dfb0882ad73acf84710
```

Runtime output

(X => 3, Y => 4, Z => 5)

In this case, only the Z component has a named association, while the other components have a positional association.

Note that a positional association cannot follow a named association, so we cannot write $P := (3, Y \Rightarrow 4, 5)$; for example. Once we start using a named association for a component, we have to continue using it for the remaining components.

In addition, we can choose multiple components at once and assign the same value to them. For that, we use the | syntax:

Listing 8: show_record_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show_Record_Aggregates is
3
      P : Point_3D;
4
  begin
5
      -- Multiple component selection
6
      P := (X | Y => 5,
7
            Ζ
                  => 6);
8
```

```
9
10 Display (P);
11 end Show_Record_Aggregates;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Pos_Named_

←Rec_Aggregates

MD5: a4fde562fb60d290caf46d86b13e694b
```

Runtime output

(X => 5, Y => 5, Z => 6)

Here, we assign 5 to both X and Y.

1 In the Ada Reference Manual

• 4.3.1 Record Aggregates¹⁰⁹

5.2.1 <>

We can use the <> syntax to tell the compiler to use the default value for specific components. However, if there's no default value for specific components, that component isn't initialized to a known value. For example:

	Listing 9:	show	record	aggreg	gates.adb
--	------------	------	--------	--------	-----------

```
with Points; use Points;
1
2
   procedure Show_Record_Aggregates is
3
      P : Point_3D;
4
   begin
5
      P := (0, 1, 2);
6
      Display (P);
7
8
      -- Specifying X component.
9
      P := (X => 42,
10
             Y => <>,
11
12
             Z => <>);
      Display (P);
13
14
      -- Specifying Y and Z components.
15
      P := (X => <>,
16
             Y => 10,
17
             Z => 20);
18
      Display (P);
19
   end Show_Record_Aggregates;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Pos_Named_

→Rec_Aggregates

MD5: 25145e7cba5a566c518ac4218e550899
```

Runtime output

¹⁰⁹ http://www.ada-auth.org/standards/22rm/html/RM-4-3-1.html

Here, as the components of Point_3D don't have a default value, those components that have <> are not initialized:

- when we write (X => 42, Y => <>, Z => <>), only X is initialized;
- when we write (X => <>, Y => 10, Z => 20) instead, only X is uninitialized.

For further reading...

As we've just seen, all components that get a <> are uninitialized because the components of Point_3D don't have a default value. As no initialization is taking place for those components of the aggregate, the actual value that is assigned to the record is undefined. In other words, the resulting behavior might dependent on the compiler's implementation.

When using GNAT, writing (X => 42, Y => <>, Z => <>) keeps the value of Y and Z intact, while (X => <>, Y => 10, Z => 20) keeps the value of X intact.

If the components of Point_3D had default values, those would have been used. For example, we may change the type declaration of Point_3D and use default values for each component:

Listing 10: points.ads

```
package Points is
1
2
      type Point 3D is record
3
         X : Integer := 10;
4
         Y : Integer := 20;
5
         Z : Integer := 30;
6
      end record;
7
8
      procedure Display (P : Point_3D);
9
10
  end Points;
11
```

Code block metadata

Then, writing <> makes use of those default values we've just specified:

Listing 11: show_record_aggregates.adb

```
with Points; use Points;
procedure Show_Record_Aggregates is
P : Point_3D := (0, 0, 0);
```

```
begin
5
      -- Using default value for
6
      -- all components
7
      P := (X => <>,
8
             Y => <>,
9
             Z => <>);
10
      Display (P);
11
  end Show_Record_Aggregates;
12
```

Code block metadata

Runtime output

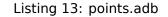
(X => 10, Y => 20, Z => 30)

Now, as expected, the default values of each component (10, 20 and 30) are used when we write <>.

Similarly, we can specify a default value for the type of each component. For example, let's declare a Point_Value type with a default value — using the Default_Value aspect — and use it in the Point_3D record type:

Listing 12: points.ads

```
package Points is
1
2
      type Point_Value is new Float
3
        with Default_Value => 99.9;
4
5
      type Point_3D is record
6
         X : Point Value;
7
         Y : Point_Value;
8
         Z : Point Value;
9
      end record;
10
11
      procedure Display (P : Point_3D);
12
13
   end Points;
14
```



```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Points is
3
4
      procedure Display (P : Point_3D) is
5
      begin
6
          Put Line ("(X => "
7
                    & Point_Value'Image (P.X)
8
                    & ",");
9
          Put_Line (" Y => "
10
                    & Point_Value'Image (P.Y)
11
                    & ",");
12
         Put_Line (" Z => "
13
                    & Point_Value'Image (P.Z)
14
```

```
15 & δ<sub>4</sub> ")");
end Display;
17
18 end Points;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregate_Default_Value

MD5: 508d7f5e7d02da1677485f7d588847f6
```

Then, writing <> makes use of the default value of the Point_Value type:

```
Listing 14: show record aggregates.adb
```

```
with Points; use Points;
1
2
   procedure Show_Record_Aggregates is
3
      P : Point 3D := (0.0, 0.0, 0.0);
4
   begin
5
          Using default value of Point Value
6
      - -
      -- for all components
7
      P := (X => <>,
8
            Y => <>,
9
             Z => <>);
10
      Display (P);
11
  end Show Record Aggregates;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregate_Default_Value

MD5: 895799077af4a295c250480c32954a2c
```

Runtime output

(X => 9.99000E+01, Y => 9.99000E+01, Z => 9.99000E+01)

In this case, the default value of the Point_Value type (99.9) is used for all components when we write <>.

5.2.2 others

Also, we can use the **others** selector to assign a value to all components that aren't explicitly mentioned in the aggregate. For example:

Listing 15: show_record_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show Record Aggregates is
3
      P : Point 3D;
4
   begin
5
      -- Specifying X component;
6
      -- using 42 for all
7
      -- other components.
8
      P := (X
               => 42,
9
            others => 100);
10
```

```
Display (P);
Display (P);
Picture Contents
Picture Contents => 256);
Display (P);
end Show_Record_Aggregates;
```

Code block metadata

Runtime output

When we write P := $(X \Rightarrow 42, \text{ others } \Rightarrow 100)$, we're assigning 42 to X and 100 to all other components (Y and Z in this case). Also, when we write P := $(\text{others } \Rightarrow 256)$, all components have the same value (256).

Note that writing a specific value in **others** — such as (**others** => 256) — only works when all components have the same type. In this example, all components of Point_3D have the same type: **Integer**. If we had components with different types in the components selected by **others**, say **Integer** and **Float**, then (**others** => 256) would trigger a compilation error. For example, consider this package:

Listing 16: custom_records.ads

```
1 package Custom_Records is
2
3 type Integer_Float is record
4 A, B : Integer := 0;
5 Y, Z : Float := 0.0;
6 end record;
7
8 end Custom_Records;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregates_Others

MD5: 875e470aa2cbc5fcfefae649ed5528f6
```

If we had written an aggregate such as (**others** => 256) for an object of type Integer_Float, the value (256) would be OK for components A and B, but not for components Y and Z:

Listing 17: show_record_aggregates_others.adb

```
with Custom_Records; use Custom_Records;
procedure Show_Record_Aggregates_Others is
Dummy : Integer_Float;
begin
    -- ERROR: components selected by
```

```
7 -- others must be of same
8 -- type.
9 Dummy := (others => 256);
10 end Show_Record_Aggregates_Others;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⊶Aggregates_Others

MD5: d543ee07e24caf63384ab0d140054be2
```

Build output

We can fix this compilation error by making sure that **others** only refers to components of the same type:

Listing 18: show_record_aggregates_others.adb

```
with Custom Records; use Custom Records;
1
2
   procedure Show_Record_Aggregates_Others is
3
      Dummy : Integer_Float;
4
   beain
5
         OK: components selected by
      - -
6
              others have Integer type.
      - -
7
      Dummy := (Y | Z => 256.0,
8
                 others => 256);
9
   end Show Record Aggregates Others;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregates_Others

MD5: d01977a49e08d2c6cb6b7788581ed56f
```

In any case, writing (**others** => <>) is always accepted by the compiler because it simply selects the default value of each component, so the type of those values is unambiguous:

Listing 19: show_record_aggregates_others.adb

```
with Custom_Records; use Custom_Records;
procedure Show_Record_Aggregates_Others is
Dummy : Integer_Float;
begin
Dummy := (others => <>);
end Show Record Aggregates Others;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregates_Others

MD5: db9b72ffc933436e76305887276eeafd
```

This code compiles because <> uses the appropriate default value of each component.

5.2.3 Record discriminants

When a record type has discriminants, they must appear as components of an aggregate of that type. For example, consider this package:

```
Listing 20: points.ads
```

```
package Points is
1
2
      type Point_Dimension is (Dim_1, Dim_2, Dim_3);
3
4
       type Point (D : Point Dimension) is record
5
          case D is
6
          when Dim_1 =>
7
             X1
                         : Integer;
8
          when Dim 2 =>
9
             X2, Y2
                         : Integer;
10
          when Dim 3 =>
11
             X3, Y3, Z3 : Integer;
12
          end case;
13
      end record;
14
15
      procedure Display (P : Point);
16
17
   end Points;
18
```

Listing 21: points.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Points is
3
4
       procedure Display (P : Point) is
5
       begin
6
          Put_Line (Point_Dimension'Image (P.D));
7
8
          case P.D is
9
          when Dim 1 =>
10
             Put_Line (" (X => "
11
                        & Integer'Image (P.X1)
12
                        & ")");
13
          when Dim_2 =>
14
             Put_Line (" (X => "
15
                        & Integer'Image (P.X2)
16
                        & ",");
17
             Put_Line ("
                             Y =>
18
                        & Integer'Image (P.Y2)
19
                       & ")");
20
          when Dim_3 =>
21
             Put_Line (" (X => "
22
                        & Integer'Image (P.X3)
23
                        & ",");
("Y=>"
24
             Put_Line ("
25
                        & Integer'Image (P.Y3)
26
                        & ",");
(" Z => "
27
             Put_Line ("
28
                        & Integer'Image (P.Z3)
29
                        & ")");
30
          end case;
31
       end Display;
32
33
   end Points;
34
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregate_Discriminant

MD5: bd71322a65ca50e1eefa0aedd407931a
```

To write aggregates of the Point type, we have to specify the D discriminant as a component of the aggregate. The discriminant must be included in the aggregate — and must be static — because the compiler must be able to examine the aggregate to determine if it is both complete and consistent. All components must be accounted for one way or another, as usual — but, in addition, references to those components whose existence depends on the discriminant's values must be consistent with the actual discriminant value used in the aggregate. For example, for type Point, an aggregate can only reference the X3, Y3, and Z3 components when Dim_3 is specified for the discriminant D; otherwise, those three components don't exist in that aggregate. Also, the discriminant D must be the first one if we use positional component association. For example:

Listing 22: show_rec_aggregate_discriminant.adb

```
with Points; use Points;
1
2
   procedure Show Rec Aggregate Discriminant is
3
       -- Positional component association
4
      P1 : constant Point := (Dim_1, 0);
5
6
       -- Named component association
7
      P2 : constant Point := (D => Dim 2,
8
                                X2 => 3.
9
                                Y2 => 4);
10
11
       -- Positional / named component association
12
      P3 : constant Point := (Dim 3,
13
                                X3 => 3,
14
                                Y3 => 4,
15
                                Z3 => 5);
16
   begin
17
      Display (P1);
18
      Display (P2);
19
      Display (P3);
20
   end Show Rec Aggregate Discriminant;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Record_Aggregates.Rec_

⇔Aggregate_Discriminant

MD5: d487e0c68ea69c3e0f2adb8ac958e31d
```

Runtime output

DIM_1 (X => 0) DIM_2 (X => 3, Y => 4) DIM_3 (X => 3, Y => 4, Z => 5)

As we see in this example, we can use any component association in the aggregate, as long as we make sure that the discriminants of the type appear as components — and are the first components in the case of positional component association.

5.3 Full coverage rules for Aggregates

\rm Note

This section was originally written by Robert A. Duff and published as Gem #1: Limited Types in Ada 2005¹¹⁰.

One interesting feature of Ada are the *full coverage rules* for aggregates. For example, suppose we have a record type:

Listing 23: persons.ads

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   package Persons is
4
      type Years is new Natural;
5
6
      type Person is record
7
         Name : Unbounded String;
8
         Age : Years;
9
      end record;
10
  end Persons;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Full_Coverage_Rules_Aggregates.

→Full_Coverage_Rules

MD5: 7755bffa8b4473c425ae5075e9c478e9
```

We can create an object of the type using an aggregate:

Listing 24: show_aggregate_init.adb

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   with Persons; use Persons;
4
5
   procedure Show_Aggregate_Init is
6
7
      X : constant Person :=
8
             (Name =>
9
                To Unbounded String ("John Doe"),
10
              Age => 25);
11
   begin
12
      null;
13
   end Show_Aggregate_Init;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Full_Coverage_Rules_Aggregates.

→Full_Coverage_Rules

MD5: 681e665b76265eff4c4d870ec011ba37
```

The full coverage rules say that every component of Person must be accounted for in the aggregate. If we later modify type Person by adding a component:

¹¹⁰ https://www.adacore.com/gems/gem-1

Listing 25: persons.ads

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   package Persons is
4
      type Years is new Natural;
5
6
      type Person is record
7
         Name
                   : Unbounded String;
8
         Age
                    : Natural;
9
         Shoe_Size : Positive;
10
      end record;
11
   end Persons;
12
```

Code block metadata

and we forget to modify X accordingly, the compiler will remind us. Case statements also have full coverage rules, which serve a similar purpose.

Of course, we can defeat the full coverage rules by using **others** (usually for *array aggregates* (page 262) and case statements, but occasionally useful for *record aggregates* (page 249)):

Listing 26: show_aggregate_init_others.adb

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   with Persons; use Persons;
4
5
   procedure Show_Aggregate_Init_Others is
6
7
      X : constant Person :=
8
             (Name
                     =>
9
                To Unbounded String ("John Doe"),
10
              others => 25);
11
   begin
12
      null;
13
  end Show_Aggregate_Init_Others;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Full_Coverage_Rules_Aggregates.

→Full_Coverage_Rules

MD5: 6d26de8dd6820682cb9150dcbb40f106
```

According to the Ada RM, **others** here means precisely the same thing as Age | Shoe_Size. But that's wrong: what **others** really means is "all the other components, including the ones we might add next week or next year". That means you shouldn't use **others** unless you're pretty sure it should apply to all the cases that haven't been invented yet.

Later on, we'll discuss full coverage rules for limited types (page 816).

5.4 Array aggregates

We've already discussed array aggregates in the Introduction to Ada¹¹¹ course. Therefore, this section just presents some details about this topic.

In the Ada Reference Manual

• 4.3.3 Array Aggregates¹¹²

5.4.1 Positional and named array aggregates

1 Note

The array aggregate syntax using brackets (e.g.: [1, 2, 3]), which we mention in this section, was introduced in Ada 2022.

Similar to *record aggregates* (page 249), array aggregates can be positional or named. Consider this package:

Listing 27: points.ads

```
package Points is
type Point_3D is array (1 .. 3) of Integer;
procedure Display (P : Point_3D);
end Points;
```

Listing 28: points.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Points is
3
4
       procedure Display (P : Point_3D) is
5
       begin
6
          Put Line ("(X => "
7
                     & Integer'Image (P (1))
8
          & ",");
Put_Line (" Y => "
9
10
                     & Integer'Image (P (2))
11
                     & ",");
12
          Put_Line (" Z => "
13
                     & Integer'Image (P (3))
14
                     & ")");
15
       end Display;
16
17
   end Points;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_ ⇔Aggregates MD5: d4b3becacc321d20810c3c90f4d8b7ff

¹¹¹ https://learn.adacore.com/courses/intro-to-ada/chapters/arrays.html#intro-ada-array-type-declaration
¹¹² http://www.ada-auth.org/standards/22rm/html/RM-4-3-3.html

We can write positional or named aggregates when assigning to an object P of Point_3D type:

Listing 29: show_array_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show Array Aggregates is
3
      P : Point 3D;
4
   begin
5
          Positional component association
6
       - -
      P := [0, 1, 2];
7
8
      Display (P);
9
10
       -- Named component association
11
      P := [1 => 3,
12
             2 => 4,
13
             3 => 51;
14
15
      Display (P);
16
   end Show_Array_Aggregates;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_ ⇔Aggregates MD5: 2d65c026639d990e7f6a99f7616d7eb4

Runtime output

In this example, we assign a positional array aggregate ([1, 2, 3]) to P. Then, we assign a named array aggregate ([1 => 3, 2 => 4, 3 => 5]) to P. In this case, the *names* are the indices of the components we're assigning to.

We can also assign array aggregates to slices:

```
Listing 30: show_array_aggregates.adb
```

```
with Points; use Points;
1
2
   procedure Show Array Aggregates is
3
      P : Point 3D := [others => 0];
4
   begin
5
          Positional component association
6
      P (2 .. 3) := [1, 2];
7
8
      Display (P);
9
10
       -- Named component association
11
      P (2 .. 3) := [1 => 3,
12
                       2 \implies 4];
13
14
      Display (P);
15
   end Show_Array_Aggregates;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates

MD5: d4e4d3ab4b7d538fc4ef1e92d28e47d9
```

Runtime output

Note that, when using a named array aggregate, the index (*name*) that we use in the aggregate doesn't have to match the slice. In this example, we're assigning the component from index 1 of the aggregate to the component of index 2 of the array P (and so on).

```
• Historically
In the first versions of Ada, we could only write array aggregates using parentheses.
                      Listing 31: show_array_aggregates.adb
 with Points; use Points;
 2
    procedure Show_Array_Aggregates is
 3
       P : Point 3D;
 4
    begin
 5

    Positional component association

 6
       P := (0, 1, 2);
 7
 8
       Display (P);
 9
 10
       -- Named component association
 11
       P := (1 => 3)
 12
             2 => 4.
 13
              3 => 5);
 14
 15
       Display (P);
 16
 17 end Show Array Aggregates;
    Code block metadata
    Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.
     →Array_Aggregates
    MD5: 16df9c01e46623ca735b84167a11a0fd
    Runtime output
     (X => 0,
     Y => 1,
     Z => 2)
     (X => 3,
     Y => 4,
     Z => 5)
    This syntax is considered obsolescent since Ada 2022: brackets ([1, 2, 3])
    should be used instead.
```

5.4.2 Null array aggregate

\rm Note

This feature was introduced in Ada 2022.

We can also write null array aggregates: []. As the name implies, this kind of array aggregate doesn't have any components.

Consider this package:

```
Listing 32: integer_arrays.ads
```

```
1 package Integer_Arrays is
2
3 type Integer_Array is
4 array (Positive range <>) of Integer;
5
6 procedure Display (A : Integer_Array);
7
8 end Integer_Arrays;
```

Listing 33: integer_arrays.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Integer_Arrays is
3
4
       procedure Display (A : Integer_Array) is
5
      begin
6
          Put_Line ("Length = "
7
                     & A'Length'Image);
8
9
          Put Line ("(");
10
          for I in A'Range loop
11
             Put ("
                      - 11
12
                   & I'Image
13
                   & " => "
14
                   & A (I)'Image);
15
             if I /= A'Last then
16
                 Put_Line (",");
17
18
             else
                 New_Line;
19
             end if;
20
          end loop;
21
          Put Line (")");
22
      end Display;
23
24
   end Integer Arrays;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates_2

MD5: 8e6e4951c14dcc6e8dea9b6a76064930
```

We can initialize an object N of Integer_Array type with a null array aggregate:

Listing 34: show_array_aggregates.adb

```
with Integer_Arrays; use Integer_Arrays;
procedure Show_Array_Aggregates is
N : constant Integer_Array := [];
begin
Display (N);
end Show_Array_Aggregates;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates_2

MD5: 188f7b006c08927f8cad83557a5e1cd9
```

Runtime output

```
Length = 0
(
)
```

In this example, when we call the Display procedure, we confirm that N doesn't have any components.

5.4.3 |, <>, others

We've seen the following syntactic elements when we were discussing *record aggregates* (page 249): |, <> and **others**. We can apply them to array aggregates as well:

Listing 35: show_array_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show_Array_Aggregates is
3
      P : Point_3D;
4
   begin
5
          All components have a value of zero.
6
      P := [others => 0];
7
8
      Display (P);
9
10
       -- Both first and second components have
11
       -- a value of three.
12
      P := [1 | 2 => 3,
13
             3
                   => 4];
14
15
      Display (P);
16
17
       -- The default value is used for the first
18
      component, and all other componentshave a value of five.
19
20
      P := [1
                => <>,
21
             others => 5];
22
23
      Display (P);
24
   end Show_Array_Aggregates;
25
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_ ⇔Aggregates MD5: 648d68f393107b138c6390c599d3d247

Runtime output

In this example, we use the |, <> and **others** elements in a very similar way as we did with record aggregates. (See the comments in the code example for more details.)

Note that, as for record aggregates, the <> makes use of the default value (if it is available). We discuss this topic in more details *later on* (page 276).

5.4.4 ...

We can also use the range syntax (...) with array aggregates:

```
Listing 36: show_array_aggregates.adb
```

```
with Points; use Points;
1
2
   procedure Show_Array_Aggregates is
3
      P : Point_3D;
4
   begin
5
      -- All components have a value of zero.
6
      P := [1 ... 3 => 0];
7
8
      Display (P);
9
10
       -- Both first and second components have
11
      -- a value of three.
12
      P := [1 .. 2 => 3,
13
             3
                    => 4];
14
15
      Display (P);
16
17
       -- The default value is used for the first
18
      component, and all other componentshave a value of five.
19
20
      P := [1
                => <>,
21
             2 ... 3 => 5];
22
23
      Display (P);
24
   end Show_Array_Aggregates;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_
Aggregates
MD5: 656f44d37ce676b24e9d512639fd0adc
```

Runtime output

(X => 0, Y => 0, Z => 0) (X => 3, Y => 3, Z => 4) (X => 1012241168, Y => 5, Z => 5)

This example is a variation of the previous one. However, in this case, we're using ranges instead of the | and **others** syntax.

5.4.5 Missing components

All aggregate components must have an associated value. If we don't specify a value for a certain component, an exception is raised:

Listing 37: show_array_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show Array Aggregates is
3
      P : Point 3D;
4
   begin
5
      P := [1 \Rightarrow 4];
6
       -- ERROR: value of components at indices
7
                  2 and 3 are missing
8
9
      Display (P);
10
  end Show Array Aggregates;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

Aggregates

MD5: 34bbc8e8bd0bd3f8b63d07fa881233bd
```

Build output

Runtime output

```
raised CONSTRAINT_ERROR : show_array_aggregates.adb:6 length check failed
```

We can use **others** to specify a value to all components that haven't been explicitly mentioned in the aggregate:

Listing 38: show_array_aggregates.adb

```
with Points; use Points;
procedure Show_Array_Aggregates is
P : Point_3D;
```

```
s begin
6 P := [1 => 4, others => 0];
7 -- 0K: unspecified components have a
8 -- value of zero
9
10 Display (P);
11 end Show_Array_Aggregates;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates

MD5: 5f1b7e3778b7d5ec990fba9558495758
```

Runtime output

(X => 4, Y => 0, Z => 0)

However, **others** can only be used when the range is known — compilation fails otherwise:

Listing 39: show_array_aggregates.adb

```
with Integer_Arrays; use Integer_Arrays;
1
2
   procedure Show Array Aggregates is
3
      N1 : Integer Array := [others => 0];
4
      -- ERROR: range is unknown
5
6
      N2 : Integer_Array (1 .. 3) := [others => 0];
7
      -- OK: range is known
8
   begin
9
      Display (N1);
10
      Display (N2);
11
  end Show_Array_Aggregates;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates_2

MD5: e185c823ca68e9193a0b12270ffebe61
```

Build output

Of course, we could fix the declaration of N1 by specifying a range — e.g. N1 : Integer_Array $(1 \dots 10) := [others => 0];$

5.4.6 Iterated component association

\rm Note

This feature was introduced in Ada 2022.

We can use an iterated component association to specify an aggregate. This is the general syntax:

```
-- All components have a value of zero
P := [for I in 1 .. 3 => 0];
```

Let's see a complete example:

Listing 40: show_array_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show_Array_Aggregates is
3
      P : Point_3D;
4
   begin
5
       -- All components have a value of zero
6
      P := [for I in 1 .. 3 => 0];
7
8
      Display (P);
9
10
       -- Both first and second components have
11
       -- a value of three
12
      P := [for I in 1 .. 3 =>
13
               (if I = 1 or I = 2
14
                then 3
15
                else 4)];
16
17
      Display (P);
18
19
       -- The first component has a value of 99
20
       -- and all other components have a value
21
       -- that corresponds to its index
22
      P := [1 => 99,
23
             for I in 2 .. 3 => I];
24
25
      Display (P);
26
   end Show_Array_Aggregates;
27
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_
⇔Aggregates
MD5: 68bddcec76f8431b16d1c090b74c2500
```

Runtime output

In this example, we use iterated component associations in different ways:

- 1. We write a simple iteration ([for I in 1 \dots 3 => 0]).
- 2. We use a conditional expression in the iteration: [for I in 1 .. 3 => (if I = 1 or I = 2 then 3 else 4)].

3. We use a named association for the first element, and then iterated component association for the remaining components: [1 => 99, for I in 2 . . 3 => I].

So far, we've used a discrete choice list (in the **for** I **in Range** form) in the iterated component association. We could use an iterator (in the **for** E **of** form) instead. For example:

Listing 41: show_array_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show Array Aggregates is
3
      P : Point_3D := [for I in Point_3D'Range => I];
4
   begin
5
      -- Each component is doubled
6
      P := [for E of P => E * 2];
7
8
      Display (P);
9
10
      -- Each component is increased
11
       -- by one
12
      P := [for E of P => E + 1];
13
14
      Display (P);
15
  end Show_Array_Aggregates;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates

MD5: 932ebc6e51c2146a726bad68b7f2cad0
```

Runtime output

In this example, we use iterators in different ways:

1. We write [for E of P => E * 2] to double the value of each component.

2. We write [for E of P => E + 1] to increase the value of each component by one.

Of course, we could write more complex operations on E in the iterators.

5.4.7 Multidimensional array aggregates

So far, we've discussed one-dimensional array aggregates. We can also use the same constructs when dealing with multidimensional arrays. Consider, for example, this package:

Listing 42: matrices.ads

```
1 package Matrices is
2
3 type Matrix is array (Positive range <>,
4
4
5
6
7 procedure Display (M : Matrix);
```

```
9 end Matrices;
```

8

Listing 43: matrices.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Matrices is
3
4
      procedure Display (M : Matrix) is
5
6
          procedure Display_Row (M : Matrix;
7
                                    I : Integer) is
8
          begin
9
             Put Line (" (");
10
             for J in M'Range (2) loop
11
                Put ("
12
                      & J'Image
13
                      & " =>
14
                      & M (I, J)'Image);
15
                 if J /= M'Last (2) then
16
                    Put_Line (",");
17
                 else
18
                    New_Line;
19
                 end if;
20
             end loop;
21
             Put (" )");
22
          end Display Row;
23
24
      begin
25
          Put_Line ("Length (1) = "
26
                     & M'Length (1)'Image);
27
          Put_Line ("Length (2) = "
28
                     & M'Length (2)'Image);
29
30
          Put Line ("(");
31
          for I in M'Range (1) loop
32
             Display_Row (M, I);
33
             if I /= M'Last (1) then
34
                 Put_Line (",");
35
             else
36
                 New_Line;
37
             end if;
38
          end loop;
39
          Put_Line (")");
40
41
      end Display;
42
43
   end Matrices;
44
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Matrix_

⇔Aggregates

MD5: 55573272f8cc0621eef7c924cfd6366a
```

We can assign multidimensional aggregates to a matrix M using positional or named component association:

```
Listing 44: show_array_aggregates.adb
```

```
with Matrices; use Matrices;
1
2
   procedure Show_Array_Aggregates is
3
      M : Matrix (1 .. 2, 1 .. 3);
4
   begin
5
         Positional component association
      - -
6
      7
8
9
      Display (M);
10
11
      -- Named component association
12
      M := [[1 => 3,
13
             2 => 4,
14
             3 => 5],
15
            [1 => 6,
16
             2 => 7,
17
             3 => 8]];
18
19
      Display (M);
20
21
   end Show_Array_Aggregates;
22
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Matrix_

_Aggregates

MD5: fe3cb6ee62422991b444c32239f72d05
```

Runtime output

```
Length (1) = 2
Length (2) = 3
(
  (
     1 => 0,
     2 =>
          1,
     3 => 2
  ),
  (
     1 => 3,
     2 => 4,
     3 => 5
  )
)
Length (1) = 2
Length (2) = 3
(
  (
     1 => 3,
     2 => 4,
     3 => 5
 ),
  (
           6,
     1 =>
     2 =>
           7,
     3 => 8
  )
)
```

The first aggregate we use in this example is [[0, 1, 2], [3, 4, 5]]. Here, [0, 1, 2] and [3, 4, 5] are subaggregates of the multidimensional aggregate. Subaggregates don't have a type themselves, but are rather just considered part of a multidimensional aggregate (which, of course, has an array type). In this sense, a subaggregate such as [0, 1, 2] is different from a one-dimensional aggregate (such as [0, 1, 2]), even though they are written in the same way.

Strings in subaggregates

In the case of matrices using characters, we can use strings in the corresponding array aggregates. Consider this package:

Listing 45: string_lists.ads

Listing 46:	string	lists.adb
-------------	--------	-----------

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body String_Lists is
3
4
      procedure Display (SL : String_List) is
5
6
          procedure Display Row (SL : String List;
7
                                   I : Integer) is
8
          begin
9
             Put (" (");
10
             for J in SL'Range (2) loop
11
                Put (SL (I, J));
12
             end loop;
13
             Put (")");
14
          end Display_Row;
15
16
      begin
17
          Put_Line ("Length (1) = "
18
                     & SL'Length (1)'Image);
19
          Put Line ("Length (2) = "
20
                     & SL'Length (2)'Image);
21
22
          Put_Line ("(");
23
          for I in SL'Range (1) loop
24
             Display_Row (SL, I);
25
             if I /= SL'Last (1) then
26
                Put Line (",");
27
             else
28
                New Line;
29
             end if;
30
          end loop;
31
          Put Line (")");
32
       end Display;
33
34
   end String Lists;
35
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.String_

⊶Aggregates

MD5: aacdbb9aa2f3b6146d8a36ca7581fd18
```

Then, when assigning to an object SL of String_List type, we can use strings in the aggregates:

Listing 47: show_array_aggregates.adb

```
with String_Lists; use String_Lists;
1
2
    procedure Show_Array_Aggregates is
3
       SL : String_List (1 .. 2, 1 .. 3);
4
    begin
5
           Positional component association
6
       - -
       SL := ["ABC",
7
                "DEF"];
8
9
       Display (SL);
10
11
       -- Named component associations
12
       SL := [[1 => 'A',
2 => 'B',
13
14
                3 => 'C'],
15
               [1 => 'D',
2 => 'E',
16
17
                3 => 'F']];
18
19
       Display (SL);
20
21
       SL := [[1 => 'X',
22
                2 => 'Y',
23
                3 => 'Z'],
24
               [others => ' ']];
25
26
       Display (SL);
27
    end Show_Array_Aggregates;
28
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.String_
⇔Aggregates
MD5: 82c28e5d8e592403d8909b8eaa1fe356
```

Runtime output

```
Length (1) = 2

Length (2) = 3

(

(ABC),

(DEF)

)

Length (1) = 2

Length (2) = 3

(

(ABC),

(DEF)

)

Length (1) = 2

Length (2) = 3
```

((XYZ), ()

In the first assignment to SL, we have the aggregate ["ABC", "DEF"], which uses strings as subaggregates. (Of course, we can use a named aggregate and assign characters to the individual components.)

5.4.8 <> and default values

As we indicated earlier, the <> syntax sets a component to its default value — if such a default value is available. If a default value isn't defined, however, the component will remain uninitialized, so that the behavior is undefined. Let's look at more complex example to illustrate this situation. Consider this package, for example:

Listing 48: points.ads

```
package Points is
1
2
      subtype Point_Value is Integer;
3
4
      type Point 3D is record
5
         X, Y, Z : Point_Value;
6
      end record;
7
8
      procedure Display (P : Point 3D);
9
10
      type Point 3D Array is
11
12
         array (Positive range <>) of Point 3D;
13
      procedure Display (PA : Point 3D Array);
14
15
   end Points;
16
```

Listing 49: points.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Points is
3
4
      procedure Display (P : Point 3D) is
5
      begin
6
                       (X => "
          Put ("
7
               & Point_Value'Image (P.X)
8
               & ",");
9
          New_Line;
10
                        Y => "
          Put ("
11
               & Point_Value'Image (P.Y)
12
               & ",");
13
          New Line;
14
                        Z => "
          Put ("
15
               & Point_Value'Image (P.Z)
16
               & ")");
17
      end Display;
18
19
      procedure Display (PA : Point_3D_Array) is
20
      begin
21
          Put Line ("(");
22
          for I in PA'Range (1) loop
23
```

```
Put_Line ("
                            - 11
24
                         & Integer'Image (I)
25
                         & " =>");
26
              Display (PA (I));
27
              if I /= PA'Last (1) then
28
                 Put_Line (",");
29
              else
30
                 New_Line;
31
              end if;
32
          end loop;
33
          Put_Line (")");
34
       end Display;
35
36
   end Points;
37
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Rec_Array_

⇔Aggregates

MD5: ffaf3745621a30362c6aadaec2c3cef2
```

Then, let's use <> for the array components:

Listing 50: show_record_aggregates.adb

```
with Points; use Points;
1
2
   procedure Show_Record_Aggregates is
3
4
       PA : Point_3D_Array (1 .. 2);
5
   begin
       PA := [ (X => 3,
6
                Y => 4,
7
                Z => 5),
8
                (X => 6,
9
                Y => 7,
10
                Z => 8) ];
11
       Display (PA);
12
13
       -- Array components are
14
       -- uninitialized.
15
       PA := [1 => <>,
16
              2 => <>];
17
       Display (PA);
18
   end Show_Record_Aggregates;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Rec_Array_

⇔Aggregates

MD5: 4575fead51e24b1a06faf4581efad112
```

Runtime output

```
 \begin{array}{c} 1 => & \\ & (X => 3, \\ & Y => 4, \\ & Z => 5), \\ 2 => & \\ & (X => 6, \\ & Y => 7, \end{array}
```

```
Z => 8)
)
(
1 =>
(X => -1444194312,
Y => 23673,
Z => -1444419927),
2 =>
(X => 23673,
Y => 1,
Z => 0)
)
```

Because the record components (of the Point_3D type) don't have default values, they remain uninitialized when we write $[1 \Rightarrow <>, 2 \Rightarrow <>]$. (In fact, you may see *garbage* in the values displayed by the Display procedure.)

When a default value is specified, it is used whenever <> is specified. For example, we could use a type that has the Default_Value aspect in its specification:

Listing 51: integer_arrays.ads

```
package Integer_Arrays is
1
2
      type Value is new Integer
3
        with Default_Value => 99;
4
5
      type Integer_Array is
6
        array (Positive range <>) of Value;
7
8
      procedure Display (A : Integer_Array);
9
10
   end Integer_Arrays;
11
```

Listing 52: show_array_aggregates.adb

```
with Integer_Arrays; use Integer_Arrays;
1
2
   procedure Show_Array_Aggregates is
3
      N : Integer_Array (1 .. 4);
4
   begin
5
      N := [for I in N'Range => Value (I)];
6
      Display (N);
7
8
      N := [others => <>];
9
      Display (N);
10
   end Show_Array_Aggregates;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

→Aggregates_2

MD5: 8007fb4af578397d1f07ad85e09ab354
```

Runtime output

```
Length = 4
(
1 => 1,
2 => 2,
3 => 3,
```

```
4 => 4
)
Length = 4
(
1 => 99,
2 => 99,
3 => 99,
4 => 99
)
```

1 2

3 4

5

6

7 8

9 10

11

1 2

3

4

5

6 7

8

9

10

11

When writing an aggregate for the Point_3D type, any component that has <> gets the default value of the Point type (99):

```
I For further reading...
Similarly, we could specify the Default_Component_Value aspect (which we discussed
earlier on (page 66)) in the declaration of the array type:
                           Listing 53: integer arrays.ads
package Integer_Arrays is
   type Value is new Integer;
   type Integer_Array is
     array (Positive range <>) of Value
       with Default_Component_Value => 9999;
   procedure Display (A : Integer_Array);
end Integer_Arrays;
                      Listing 54: show array aggregates.adb
with Integer_Arrays; use Integer_Arrays;
procedure Show_Array_Aggregates is
   N : Integer_Array (1 .. 4);
begin
   N := [for I in N'Range => Value (I)];
   Display (N);
   N := [others => <>];
   Display (N);
end Show_Array_Aggregates;
Code block metadata
```

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_

⇔Aggregates_2

MD5: 3f535bc5ce7f74ab0f0f48098a82c98a
```

Runtime output

```
Length = 4
(
1 => 1,
2 => 2,
3 => 3,
4 => 4
)
Length = 4
(
1 => 9999,
2 => 9999,
5.4. 3Array syggregates
4 => 9999
```

)

In this case, when writing <> for a component, the value specified in the Default Component Value aspect is used.

Finally, we might want to use both Default_Value (which we discussed *previously* (page 65)) and Default_Component_Value aspects at the same time. In this case, the value specified in the Default_Component_Value aspect has higher priority:

Listing 55: integer_arrays.ads

```
package Integer_Arrays is
```

```
type Value is new Integer
with Default_Value => 99;
type Integer_Array is
```

```
array (Positive range <>) of Value
with Default_Component_Value => 9999;
```

```
procedure Display (A : Integer_Array);
```

```
end Integer_Arrays;
```

Listing 56: show_array_aggregates.adb

```
with Integer_Arrays; use Integer_Arrays;
procedure Show_Array_Aggregates is
    N : Integer_Array (1 .. 4);
begin
    N := [for I in N'Range => Value (I)];
    Display (N);
    N := [others => <>];
    Display (N);
```

```
9
10
11
```

1 2

3

4 5

6

7

8 9

10 11

12

1 2

3

4

5

6

7 8

Code block metadata

end Show_Array_Aggregates;

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Array_Aggregates.Array_
Aggregates_2
MD5: e58618b565874acaa99c5d494c2acaa4
```

Runtime output

```
Length = 4
(

1 => 1,

2 => 2,

3 => 3,

4 => 4
)

Length = 4
(

1 => 9999,

2 => 9999,

3 => 9999,

4 => 9999
)
```

Here, 9999 is used when we specify <> for a component.

5.5 Extension Aggregates

Extension aggregates provide a convenient way to express an aggregate for a type that extends — adds components to — some existing type (the "ancestor"). Although mainly a matter of convenience, an extension aggregate is essential when we want to express an aggregate for an extension of a private ancestor type, that is, when we don't have compile-time visibility to the ancestor type's components.

```
1 In the Ada Reference Manual
```

```
• 4.3.2 Extension Aggregates<sup>113</sup>
```

5.5.1 Assignments to objects of derived types

Before we discuss extension aggregates in more detail, though, let's start with a simple use-case. Let's say we have:

- an object A of tagged type T1, and
- an object B of tagged type T2, which extends T1.

We can initialize object B by:

- copying the T1 specific information from A to B, and
- initializing the T2 specific components of B.

We can translate the description above to the following code:

```
A : T1;
B : T2;
begin
T1 (B) := A;
B.Extended_Component_1 := Some_Value;
-- [...]
```

Here, we use T1 (B) to select the ancestor view of object B, and we copy all the information from A to this part of B. Then, we initialize the remaining components of B. We'll elaborate on this kind of assignments later on.

5.5.2 Example: Points

To present a more concrete example, let's start with a package that defines one, two and three-dimensional point types:

Listing 57: points.ads

```
package Points is
1
2
      type Point_1D is tagged record
3
         X : Float;
4
      end record;
5
6
      procedure Display (P : Point 1D);
7
8
      type Point 2D is new Point 1D with record
9
         Y : Float;
10
```

(continues on next page)

¹¹³ http://www.ada-auth.org/standards/22rm/html/RM-4-3-2.html

```
end record;
11
12
       procedure Display (P : Point_2D);
13
14
       type Point_3D is new Point_2D with record
15
          Z : Float;
16
       end record;
17
18
       procedure Display (P : Point_3D);
19
20
   end Points;
21
```

Listing 58: points.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Points is
3
4
      procedure Display (P : Point_1D) is
5
      begin
6
          Put_Line ("(X => " & P.X'Image & ")");
7
      end Display;
8
9
      procedure Display (P : Point_2D) is
10
       begin
11
          Put Line ("(X => " & P.X'Image
12
                    & ", Y => " & P.Y'Image & ")");
13
       end Display;
14
15
       procedure Display (P : Point_3D) is
16
17
      begin
          Put_Line ("(X => " & P.X'Image
18
                    & ", Y => " & P.Y'Image
19
                    & ", Z => " & P.Z'Image & ")");
20
      end Display;
21
22
   end Points;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_

⇔Aggregate_Points

MD5: 0acc05ae2310ab4ba038dfdb6bae0495
```

Let's now focus on the Show_Points procedure below, where we initialize a two-dimensional point using a one-dimensional point.

Listing 59: show points.adb

```
with Points; use Points;
1
2
   procedure Show Points is
3
       P 1D : Point 1D;
4
       P_2D : Point_2D;
5
   begin
6
       P_{1D} := (X \implies 0.5);
7
       Display (P_1D);
8
9
       Point_1D (P_2D) := P_1D;
10
       -- Equivalent to: "P_2D.X := P_1D.X;"
11
12
```

13 P_2D.Y := 0.7; 14 15 Display (P_2D); 16 end Show_Points;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_

⇔Aggregate_Points

MD5: 68ae6fa8e6f779aebea97085bd75e082
```

Runtime output

```
(X => 5.00000E-01)
(X => 5.00000E-01, Y => 7.00000E-01)
```

In this example, we're initializing P_2D using the information stored in P_1D. By writing Point_1D (P_2D) on the left side of the assignment, we specify that we want to limit our focus on the Point_1D view of the P_2D object. Then, we assign P_1D to the Point_1D view of the P_2D object. Then, we assign P_1D to the Point_1D view of the P_2D object. The Point_2D specific components are not changed by this assignment. (In other words, this is equivalent to just writing P_2D.X := P_1D.X, as the Point_1D type only has the X component.) Finally, in the next line, we initialize the Y component with 0.7.

5.5.3 Using extension aggregates

Note that, in the assignment to P_1D, we use a record aggregate. Extension aggregates are similar to record aggregates, but they include the **with** keyword — for example: (0bj1 **with** Y => 0.5). This allows us to assign to an object with information from another object 0bj1 of a parent type and, in the same expression, set the value of the Y component of the type extension.

Let's rewrite the previous Show_Points procedure using extension aggregates:

Listing 60: show_points.adb

```
with Points; use Points;
1
2
   procedure Show Points is
З
      P_1D : Point 1D;
4
      P 2D : Point 2D;
5
   begin
6
       P_{1D} := (X => 0.5);
7
      Display (P 1D);
8
9
       P 2D := (P 1D with Y => 0.7);
10
      Display (P_2D);
11
   end Show_Points;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_

⇔Aggregate_Points

MD5: 4d03f6a565126b602d6f21fe5ee6dd27
```

Runtime output

```
(X => 5.00000E-01)
(X => 5.00000E-01, Y => 7.00000E-01)
```

When we write $P_{2D} := (P_{1D} \text{ with } Y \Rightarrow 0.7)$, we're initializing P_{2D} using:

- the information from the P_1D object of Point_1D type, which is an ancestor of the Point_2D type —, and
- the information from the record component association list for the remaining components of the Point_2D type. (In this case, the only remaining component of the Point_2D type is Y.)

We could also specify the type of the extension aggregate. For example, in the previous assignment to P_2D, we could write Point_2D'(...) to indicate that we expect the Point_2D type for the extension aggregate.

```
-- Explicitly state that the type of the
-- extension aggregate is Point_2D:
P_2D := Point_2D'(P_1D with Y => 0.7);
```

Also, we don't have to use named association in extension aggregates. We could just use positional association instead. Therefore, we could simplify the assignment to P_2D in the previous example by just writing:

P_2D := (P_1D with 0.7);

5.5.4 More extension aggregates

We can use extension aggregates for descendants of the Point_2D type as well. For example, let's extend our previous code example by declaring an object of Point_3D type (called P_3D) and use extension aggregates in assignments to this object:

```
with Points; use Points;
1
2
   procedure Show_Points is
3
       P 1D : Point 1D;
4
       P 2D : Point 2D;
5
       P 3D : Point 3D;
6
   begin
7
       P_{1D} := (X \implies 0.5);
8
       Display (P_1D);
9
10
       P_2D := (P_1D with Y => 0.7);
11
       Display (P_2D);
12
13
       P_3D := (P_2D with Z => 0.3);
14
       Display (P_3D);
15
16
       P_{3D} := (P_{1D} \text{ with } Y | Z => 0.1);
17
       Display (P_3D);
18
   end Show_Points;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_

⇔Aggregate_Points

MD5: 2ec6831557c43f697bffce8496962b53
```

Runtime output

```
(X => 5.00000E-01)
(X => 5.00000E-01, Y => 7.00000E-01)
(X => 5.00000E-01, Y => 7.00000E-01, Z => 3.00000E-01)
(X => 5.00000E-01, Y => 1.00000E-01, Z => 1.00000E-01)
```

In the first assignment to P_3D in the example above, we're initializing this object with information from P_2D and specifying the value of the Z component. Then, in the next assignment to the P_3D object, we're using an aggregate with information from P_1 and specifying values for the Y and Z components. (Just as a reminder, we can write Y $\mid Z => 0.1$ to assign 0.1 to both Y and Z components.)

5.5.5 with others

Other versions of extension aggregates are possible as well. For example, we can combine keywords and write with others to focus on all remaining components of an extension aggregate.

Listing 62: show	points.adb
------------------	------------

```
with Points; use Points;
1
2
   procedure Show Points is
3
       P_1D : Point_1D;
4
       P_2D : Point_2D;
5
       P_3D : Point_3D;
6
7
   begin
       P_{1D} := (X \implies 0.5);
8
       P_2D := (P_1D \text{ with } Y => 0.7);
9
10
          Initialize P 3D with P 1D and set other
11
       - -
          components to 0.6.
12
13
       P 3D := (P 1D with others => 0.6);
14
       Display (P_3D);
15
16
       -- Initialize P_3D with P_2D, and other
17
       -- components with their default value.
18
19
       P 3D := (P 2D with others => <>);
20
       Display (P 3D);
21
   end Show_Points;
22
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_

⇔Aggregate_Points

MD5: 0594586fc59ead106258cef8682927e9
```

Runtime output

```
(X => 5.00000E-01, Y => 6.00000E-01, Z => 6.00000E-01)
(X => 5.00000E-01, Y => 7.00000E-01, Z => 2.59244E+14)
```

In this example, the first assignment to P_3D has an aggregate with information from P_1D , while the remaining components — in this case, Y and Z — are just set to 0.6.

Continuing with this example, in the next assignment to P_3D, we're using information from P_2 in the extension aggregate. This covers the Point_2D part of the P_3D object — components X and Y, to be more specific. The Point_3D specific components of P_3D — component Z in this case — receive their corresponding default value. In this specific case, however, we haven't specified a default value for component Z in the declaration of the Point_3D type, so we cannot rely on any specific value being assigned to that component when using **others** => <>.

5.5.6 with null record

We can also use extension aggregates with null records. Let's focus on the P_3D_Ext object of Point_3D_Ext type. This object is declared in the Show_Points procedure of the next code example.

```
Listing 63: points-extensions.ads
```

```
package Points.Extensions is

type Point_3D_Ext is new
Point_3D with null record;

end Points.Extensions;
```

Listing 64: show points.adb

```
with Points;
                              use Points;
1
   with Points.Extensions; use Points.Extensions;
2
3
   procedure Show Points is
4
      P 3D
                : Point_3D;
5
      P_3D_Ext : Point_3D_Ext;
6
   begin
7
      P_3D := (X \implies 0.0, Y \implies 0.5, Z \implies 0.4);
8
9
      P_3D_Ext := (P_3D with null record);
10
      Display (P_3D_Ext);
11
  end Show Points;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_ ⇔Aggregate_Points MD5: 8ec3ddb3a1f2a6e550ac4d622e97124c

Runtime output

(X => 0.00000E+00, Y => 5.00000E-01, Z => 4.00000E-01)

The P_3D_Ext object is of Point_3D_Ext type, which is declared in the Points.Extensions package and derived from the Point_3D type. Note that we're not extending Point_3D_Ext with new components, but using a null record instead in the declaration. Therefore, as the Point_3D_Ext type doesn't own any new components, we just write (P_3D with null record) to initialize the P_3D_Ext object.

5.5.7 Extension aggregates and descendent types

In the examples above, we've been initializing objects of descendent types by using objects of ascending types in extension aggregates. We could, however, do the opposite and initialize objects of ascending types using objects of descendent type in extension aggregates. Consider this code example:

Listing 65: show_points.adb

```
with Points; use Points;
procedure Show_Points is
P_2D : Point_2D;
P_3D : Point_3D;
begin
```

```
7 P_3D := (X => 0.5, Y => 0.7, Z => 0.3);
8 Display (P_3D);
9 P_2D := (Point_1D (P_3D) with Y => 0.3);
11 Display (P_2D);
12 end Show_Points;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Extension_Aggregates.Extension_

⇔Aggregate_Points

MD5: ae5e88a36c58b1eb495d5ba8752e50e7
```

Runtime output

```
(X => 5.00000E-01, Y => 7.00000E-01, Z => 3.00000E-01)
(X => 5.00000E-01, Y => 3.00000E-01)
```

Here, we're using Point_1D (P_3D) to select the Point_1D view of an object of Point_3D type. At this point, we have specified the Point_1D part of the aggregate, so we still have to specify the remaining components of the Point_2D type — the Y component, to be more specific. When we do that, we get the appropriate aggregate for the Point_2D type. In summary, by carefully selecting the appropriate view, we're able to initialize an object of a scending type (Point_2D), which contains less components, using an object of a descendent type (Point_3D), which contains more components.

5.6 Delta Aggregates

\rm 1 Note

This feature was introduced in Ada 2022.

Previously, we've discussed *extension aggregates* (page 283), which are used to assign an object Obj_From of a tagged type to an object Obj_To of a descendent type.

We may want also to assign an object Obj_From of to an object Obj_To of the same type, but change some of the components in this assignment. To do this, we use delta aggregates.

5.6.1 Delta Aggregates for Tagged Records

Let's reuse the Points package from a previous example:

```
Listing 66: points.ads
```

```
package Points is
1
2
      type Point 1D is tagged record
3
         X : Float;
4
      end record;
5
6
7
      type Point 2D is new Point 1D with record
         Y : Float;
8
      end record;
9
10
       type Point_3D is new Point_2D with record
11
          Z : Float;
12
```

```
13 end record;
14
15 procedure Display (P : Point_3D);
16
17 end Points;
```

Listing 67: points.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
    package body Points is
3
4
       procedure Display (P : Point_3D) is
5
       begin
6
           Put Line ("(X => " & P.X'Image
7
                      & ", Y => " & P.Y'Image
& ", Z => " & P.Z'Image & ")");
8
9
       end Display;
10
11
   end Points;
12
```

Listing 68: show points.adb

```
1
   with Points; use Points;
2
   procedure Show_Points is
3
       P1, P2, P3 : Point_3D;
4
   begin
5
       P1 := (X \implies 0.5, Y \implies 0.7, Z \implies 0.3);
6
       Display (P1);
7
8
       P2 := (P1 with delta X => 1.0);
9
       Display (P2);
10
11
       P3 := (P1 with delta X => 0.2, Y => 0.3);
12
       Display (P3);
13
   end Show_Points;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Delta_Aggregates.Delta_

⇔Aggregates_Tagged

MD5: 23e9f53d626e32fc0524abfa0a437dbf
```

Runtime output

```
(X => 5.00000E-01, Y => 7.00000E-01, Z => 3.00000E-01)
(X => 1.00000E+00, Y => 7.00000E-01, Z => 3.00000E-01)
(X => 2.00000E-01, Y => 3.00000E-01, Z => 3.00000E-01)
```

Here, we assign P1 to P2, but change the X component. Also, we assign P1 to P3, but change the X and Y components.

We can use class-wide types with delta aggregates. Consider this example:

Listing 69: show_points.adb

```
with Points; use Points;
procedure Show_Points is
```

```
P_3D : Point_3D;
5
6
       function Reset (P_2D : Point_2D'Class)
7
                          return Point_2D'Class is
8
         ((P_2D with delta X | Y => 0.0));
9
10
   begin
11
       P_{3D} := (X \implies 0.1, Y \implies 0.2, Z \implies 0.3);
12
       Display (P_3D);
13
14
       P_3D := Point_3D (Reset (P_3D));
15
       Display (P_3D);
16
17
   end Show_Points;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Delta_Aggregates.Delta_

→Aggregates_Tagged

MD5: dca144fe420dd37e224d089458f9e8a8
```

Runtime output

```
(X => 1.00000E-01, Y => 2.00000E-01, Z => 3.00000E-01)
(X => 0.00000E+00, Y => 0.00000E+00, Z => 3.00000E-01)
```

In this example, the Reset function returns an object of Point_2D'Class where all components of Point_2D'Class type are zero. We call the Reset function for the P_3D object of Point_3D type, so that only the Z component remains untouched.

Note that we use the syntax X \mid Y in the body of the Reset function and assign the same value to both components.

For further reading...

We could have implemented Reset as a procedure — in this case, without using delta aggregates:

```
Listing 70: show points.adb
```

```
with Points; use Points;
1
2
3
    procedure Show_Points is
4
       P_3D : Point_3D;
5
6
       procedure Reset
7
         (P_2D : in out Point_2D'Class) is
8
       begin
9
           Point_2D (P_2D) := (others => 0.0);
10
       end Reset;
11
12
    begin
13
       P_3D := (X \implies 0.1, Y \implies 0.2, Z \implies 0.3);
14
15
       Display (P_3D);
16
       Reset (P_3D);
17
       Display (P_3D);
18
19
   end Show Points;
20
```

5.6.2 Delta Aggregates for Non-Tagged Records

The examples above use tagged types. We can also use delta aggregates with non-tagged types. Let's rewrite the Points package and convert Point_3D to a non-tagged record type.

Listing 71: points.ads

```
package Points is
1
2
       type Point 3D is record
3
          X : Float;
4
          Y : Float;
5
          Z : Float;
6
      end record;
7
8
      procedure Display (P : Point_3D);
9
10
   end Points;
11
```

Listing 72: points.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
    package body Points is
3
4
       procedure Display (P : Point_3D) is
5
       begin
6
           Put_Line ("(X => " & P.X'Image
7
                      & ", Y => " & P.Y'Image
& ", Z => " & P.Z'Image & ")");
8
9
       end Display;
10
11
   end Points;
12
```

Listing 73: show_points.adb

```
with Points; use Points;
1
2
   procedure Show_Points is
3
       P1, P2, P3 : Point_3D;
4
   begin
5
       P1 := (X \implies 0.5, Y \implies 0.7, Z \implies 0.3);
6
       Display (P1);
7
8
       P2 := (P1 with delta X => 1.0);
9
       Display (P2);
10
11
       P3 := (P1 with delta X => 0.2, Y => 0.3);
12
       Display (P3);
13
   end Show_Points;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Delta_Aggregates.Delta_

⇔Aggregates_Non_Tagged

MD5: 1f12f33ac0a84919978c56d04f479e35
```

Runtime output

```
(X => 5.00000E-01, Y => 7.00000E-01, Z => 3.00000E-01)
(X => 1.00000E+00, Y => 7.00000E-01, Z => 3.00000E-01)
(X => 2.00000E-01, Y => 3.00000E-01, Z => 3.00000E-01)
```

In this example, Point_3D is a non-tagged type. Note that we haven't changed anything in the Show_Points procedure: it still works as it did with tagged types.

5.6.3 Delta Aggregates for Arrays

We can use delta aggregates for arrays. Let's change the declaration of Point_3D and use an array to represent a 3-dimensional point:

Listing 74: points.ads

```
package Points is
1
2
      type Float Array is
3
        array (Positive range <>) of Float;
4
5
      type Point_3D is new Float_Array (1 .. 3);
6
7
      procedure Display (P : Point_3D);
8
9
   end Points;
10
```

Listing 75: points.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Points is
3
4
       procedure Display (P : Point_3D) is
5
       begin
6
          Put ("(");
7
          for I in P'Range loop
8
             Put (I'Image
9
                  & " => "
10
                   & P (I) 'Image);
11
          end loop;
12
          Put Line (")");
13
       end Display;
14
15
   end Points;
16
```

Listing 76: show_points.adb

```
with Points; use Points;
1
2
   procedure Show Points is
3
      P1, P2, P3 : Point_3D;
4
   begin
5
      P1 := [0.5, 0.7, 0.3];
6
      Display (P1);
7
8
      P2 := [P1 with delta 1 => 1.0];
9
      Display (P2);
10
11
      P3 := [P1 with delta 1 => 0.2, 2 => 0.3];
12
      -- Alternatively:
13
      -- P3 := [P1 with delta 1 ... 2 => 0.2, 0.3];
14
15
      Display (P3);
16
   end Show_Points;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Aggregates.Delta_Aggregates.Delta_ ⇔Aggregates_Array MD5: 06293882e5dd020f56fbced6bc03ccf0

Runtime output

(1 => 5.00000E-01 2 => 7.00000E-01 3 => 3.00000E-01) (1 => 1.00000E+00 2 => 7.00000E-01 3 => 3.00000E-01) (1 => 2.00000E-01 2 => 3.00000E-01 3 => 3.00000E-01)

The implementation of Show_Points in this example is very similar to the version where use a record type. In this case, we:

- assign P1 to P2, but change the first component, and
- we assign P1 to P3, but change the first and second components.

Using slices

In the assignment to P3, we can either specify each component of the delta individually or use a slice: both forms are equivalent. Also, we can use slices to assign the same number to multiple components:

Listing 77: show_points.adb

```
with Points; use Points;
1
2
   procedure Show Points is
3
      P1, P3 : Point 3D;
4
   begin
5
      P1 := [0.5, 0.7, 0.3];
6
7
      Display (P1);
8
      P3 := [P1 with delta
9
               P3'First + 1 .. P3'Last => 0.0];
10
      Display (P3);
11
  end Show_Points;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Delta_Aggregates.Delta_

⇔Aggregates_Array

MD5: 0a00e17b2d803f23edc728969d663c59
```

Runtime output

```
(1 => 5.00000E-01 2 => 7.00000E-01 3 => 3.00000E-01)
(1 => 5.00000E-01 2 => 0.00000E+00 3 => 0.00000E+00)
```

In this example, we're assigning P1 to P3, but resetting all components of the array starting by the second one.

Multiple components

We can also assign multiple components or slices:

Listing 78: float arrays.ads

```
package Float_Arrays is
type Float Array is
```

```
4 array (Positive range <>) of Float;
5 procedure Display (P : Float_Array);
8 end Float_Arrays;
```

Listing 79: float_arrays.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
   package body Float_Arrays is
3
4
      procedure Display (P : Float Array) is
5
      begin
6
7
          Put ("(");
8
          for I in P'Range loop
9
             Put (I'Image
10
                  & " => "
11
                  & P (I)'Image);
12
          end loop;
13
          Put_Line (")");
14
15
      end Display;
16
17
   end Float Arrays;
18
```

Listing 80: show_multiple_delta_slices.adb

```
with Float_Arrays; use Float_Arrays;
1
2
   procedure Show Multiple Delta Slices is
3
4
5
      P1, P2 : Float_Array (1 .. 5);
6
   begin
7
      P1 := [1.0, 2.0, 3.0, 4.0, 5.0];
8
      Display (P1);
9
10
      P2 := [P1 with delta
11
                P2'First + 1 .. P2'Last - 2 => 0.0,
12
                P2'Last - 1 .. P2'Last => 0.2];
13
      Display (P2);
14
15
   end Show_Multiple_Delta_Slices;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Aggregates.Delta_Aggregates.Delta_

⇔Aggregates_Array

MD5: 37063cd1c6cd46522d8e5b0df7b5741b
```

Runtime output

```
(1 => 1.00000E+00 2 => 2.00000E+00 3 => 3.00000E+00 4 => 4.00000E+00 5 => 5.

G00000E+00)

(1 => 1.00000E+00 2 => 0.00000E+00 3 => 0.00000E+00 4 => 2.00000E-01 5 => 2.

G00000E-01)
```

In this example, we have two arrays P1 and P2 of Float_Array type. We assign P1 to P2, but change:

- the second to the last-but-two components to 0.0, and
- the last-but-one and last components to 0.2.

1 In the Ada Reference Manual

• Delta Aggregates¹¹⁴

¹¹⁴ http://www.ada-auth.org/standards/22rm/html/RM-4-3-4.html

ARRAYS

6.1 Array constraints

Array constraints are important in the declaration of an array because they define the total size of the array. In fact, arrays must always be constrained. In this section, we start our discussion with unconstrained array types, and then continue with constrained arrays and arrays types. Finally, we discuss the differences between unconstrained arrays and vectors.

1 In the Ada Reference Manual

• 3.6 Array Types¹¹⁵

6.1.1 Unconstrained array types

In the Introduction to Ada course¹¹⁶, we've seen that we can declare array types whose bounds are not fixed: in that case, the bounds are provided when creating objects of those types. For example:

Listing 1: measurement defs.ads

```
1 package Measurement_Defs is
2
3 type Measurements is
4 array (Positive range <>) of Float;
5 -- ^ Bounds are of type Positive,
6 -- but not known at this point.
7
8 end Measurement Defs;
```

Listing 2: show_measurements.adb

```
use Ada.Text I0;
   with Ada.Text I0;
1
2
   with Measurement Defs; use Measurement Defs;
3
4
   procedure Show Measurements is
5
       M : Measurements (1 .. 10);
6
                              ^ Providing bounds here!
7
   begin
8
       Put_Line ("First index: " & M'First'Image);
Put_Line ("Last index: " & M'Last'Image);
9
10
   end Show Measurements;
11
```

¹¹⁵ http://www.ada-auth.org/standards/22rm/html/RM-3-6.html
 ¹¹⁶ https://learn.adacore.com/courses/intro-to-ada/chapters/arrays.html#intro-ada-unconstrained-array-types

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Array_Constraints.Unconstrained_

⇔Array_Type

MD5: a5cdc74dd61e36476431cf675452d1d5
```

Build output

Runtime output

First index: 1 Last index: 10

In this example, the Measurements array type from the Measurement_Defs package is unconstrained. In the Show_Measurements procedure, we declare a constrained object (M) of this type.

6.1.2 Constrained arrays

The Introduction to Ada course¹¹⁷ highlights the fact that the bounds are fixed once an object is declared:

Although different instances of the same unconstrained array type can have different bounds, a specific instance has the same bounds throughout its lifetime. This allows Ada to implement unconstrained arrays efficiently; instances can be stored on the stack and do not require heap allocation as in languages like Java.

In the Show_Measurements procedure above, once we declare M, its bounds are fixed for the whole lifetime of M. We cannot *add* another component to this array. In other words, M will have 10 components for its whole lifetime:

```
M : Measurements (1 .. 10);
--
Bounds cannot be changed!
```

6.1.3 Constrained array types

Note that we could declare constrained array types. Let's rework the previous example:

```
Listing 3: measurement_defs.ads
```

```
1 package Measurement_Defs is
2
3 type Measurements is
4 array (1 .. 10) of Float;
5 -- ^ Bounds are of known and fixed.
6
7 end Measurement_Defs;
```

Listing 4: show_measurements.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Measurement_Defs; use Measurement_Defs;
```

(continues on next page)

¹¹⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/arrays.html#intro-ada-unconstrained-array-type-instance-bound

```
procedure Show_Measurements is
5
      M : Measurements;
6
                        ^ We cannot change the
7
       - -
                          bounds here!
       - -
8
   begin
9
      Put_Line ("First index: " & M'First'Image);
10
      Put_Line ("Last index: " & M'Last'Image);
11
  end Show_Measurements;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Array_Constraints.Constrained_

⇔Array_Type

MD5: 4741986fdf4dab731baa001b6e60c345
```

Build output

```
show_measurements.adb:6:04: warning: variable "M" is read but never assigned [- {\scriptstyle {\rm             _s}} gnatwv]
```

Runtime output

First index: 1 Last index: 10

In this case, the bounds of the Measurements type are fixed. Now, we cannot specify the bounds (or change them) in the declaration of the M array, as they have already been defined in the type declaration.

Unconstrained Arrays vs. Vectors

If you need, however, the flexibility of increasing the length of an array, you could use the language-defined Vector type instead. This is how we could rewrite the previous example using vectors:

Listing 5: measurement defs.ads

```
with Ada.Containers; use Ada.Containers;
1
   with Ada.Containers.Vectors;
2
3
   package Measurement_Defs is
4
5
      package Vectors is new Ada.Containers.Vectors
6
        (Index_Type => Positive,
7
         Element_Type => Float);
8
9
      subtype Measurements is Vectors.Vector;
10
11
   end Measurement_Defs;
12
```

Listing 6: show measurements.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Measurement_Defs; use Measurement_Defs;
procedure Show_Measurements is
use Measurement_Defs.Vectors;
M : Measurements := To_Vector (10);
```

```
^ Creating 10-element
9
10
                               vector.
   begin
11
      Put_Line ("First index: "
12
                  & M.First_Index'Image);
13
       Put_Line ("Last index:
14
                  & M.Last_Index'Image);
15
16
      Put_Line ("Adding element...");
17
      M.Append (1.0);
18
19
       Put_Line ("First index: "
20
                  & M.First_Index'Image);
21
       Put_Line ("Last index:
22
                  & M.Last_Index'Image);
23
   end Show_Measurements;
24
```

Code block metadata

Runtime output

First index: 1 Last index: 10 Adding element... First index: 1 Last index: 11

In the declaration of M in this example, we're creating a 10-element vector by calling To_Vector and specifying the element count. Later on, with the call to Append, we're increasing the length of the M to 11 elements.

As you might expect, the flexibility of vectors comes with a price: every time we add an element that doesn't fit in the current capacity of the vector, the container has to reallocate memory in the background due to that new element. Therefore, arrays are more efficient, as the memory allocation only happens once for each object.

In the Ada Reference Manual

- 3.6 Array Types¹¹⁸
- A.18.2 The Generic Package Containers. Vectors¹¹⁹

6.2 Multidimensional Arrays

So far, we've discussed unidimensional arrays, since they are very common in Ada. However, Ada also supports multidimensional arrays using the same facilities as for unidimensional arrays. For example, we can use the First, Last, **Range** and Length attributes for each dimension of a multidimensional array. This section presents more details on this topic.

To create a multidimensional array, we simply separate the ranges of each dimension

¹¹⁸ http://www.ada-auth.org/standards/22rm/html/RM-3-6.html

¹¹⁹ http://www.ada-auth.org/standards/22rm/html/RM-A-18-2.html

with a comma. The following example presents the one-dimensional array A1, the twodimensional array A2 and the three-dimensional array A3:

```
Listing 7: multidimensional_arrays_decl.ads
```

```
package Multidimensional Arrays Decl is
1
2
      A1 : array (1 .. 10) of Float;
3
      A2 : array (1 .. 5, 1 .. 10) of Float;
4
                    ^ first dimension
       - -
5
                            ^ second dimension
       - -
6
      A3 : array (1 .. 2, 1 .. 5, 1 .. 10) of Float;
7
                    ^ first dimension
       - -
8
                              second dimension
9
       - -
                                     ^ third dimension
10
   end Multidimensional Arrays Decl;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Multidimensional_Arrays.

→Multidimensional_Arrays

MD5: 928243b293c67a078d729c3cac68bb92
```

The two-dimensional array A2 has 5 components in the first dimension and 10 components in the second dimension. The three-dimensional array A3 has 2 components in the first dimension, 5 components in the second dimension, and 10 components in the third dimension. Note that the ranges we've selected for A1, A2 and A3 are completely arbitrary. You may select ranges for each dimension that are the most appropriate in the context of your application. Also, the number of dimensions is not limited to three, so you could declare higher-dimensional arrays if needed.

We can use the Length attribute to retrieve the length of each dimension. We use an integer value in parentheses to specify which dimension we're referring to. For example, if we write A'Length (2), we're referring to the length of the second dimension of a multidimensional array A. Note that A'Length is equivalent to A'Length (1). The same equivalence applies to other array-related attributes such as First, Last and **Range**.

Let's use the Length attribute for the arrays we declared in the Multidimensional_Arrays_Decl package:

Listing 8: show_multidimensional_arrays.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Multidimensional_Arrays_Decl;
3
   use Multidimensional Arrays Decl;
4
5
   procedure Show_Multidimensional_Arrays is
6
   begin
7
      Put_Line ("A1'Length:
8
                 & A1'Length'Image);
9
      Put_Line ("A1'Length (1):
10
                 & Al'Length (1)'Image);
11
      Put Line ("A2'Length (1):
12
                 & A2'Length (1)'Image);
13
      Put Line ("A2'Length (2):
14
                 & A2'Length (2)'Image);
15
      Put_Line ("A3'Length (1):
16
                 & A3'Length (1)'Image);
17
      Put_Line ("A3'Length (2): "
18
                 & A3'Length (2)'Image);
19
      Put Line ("A3'Length (3): "
20
```

21		& A3'Length (3)'Image);	
22	end Show	Multidimensional Arrays;	

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Arrays.Multidimensional_Arrays. →Multidimensional_Arrays MD5: 70b9b8df7e46302b92613fa484ef71ca

Runtime output

A1'Length:		10
A1'Length	(1):	10
A2'Length	(1):	5
A2'Length	(2):	10
A3'Length	(1):	2
A3'Length	(2):	5
A3'Length	(3):	10

As this simple example shows, we can easily retrieve the length of each dimension. Also, as we've just mentioned, A1'Length is equal to A1'Length (1).

Let's consider an application where we make hourly measurements for the first 12 hours of the day, on each day of the week. We can create a two-dimensional array type called Measurements to store this data. Also, we can have three procedures for this array:

- Show_Indices, which presents the indices (days and hours) of the two-dimensional array;
- Show_Values, which presents the values stored in the array; and
- Reset, which resets each value of the array.

This is the complete code for this application:

Listing 9: measurement_defs.ads

```
package Measurement_Defs is
1
2
      type Days is
3
         (Mon, Tue, Wed, Thu, Fri, Sat, Sun);
4
5
      type Hours is range 0 .. 11;
6
7
      subtype Measurement is Float;
8
9
      type Measurements is
10
        array (Days, Hours) of Measurement;
11
12
      procedure Show_Indices (M : Measurements);
13
14
      procedure Show_Values (M : Measurements);
15
16
      procedure Reset (M : out Measurements);
17
18
   end Measurement_Defs;
19
```

Listing 10: measurement_defs.adb

1	<pre>with Ada.Text_I0; use Ada.Text_I0</pre>	;
2		
3	<pre>package body Measurement_Defs is</pre>	

```
procedure Show_Indices (M : Measurements) is
5
       begin
6
          Put_Line ("---- Indices ----");
7
8
          for D in M'Range (1) loop
9
             Put (D'Image & " ");
10
11
             for H in M'First (2) ..
12
                       M'Last (2) - 1
13
             loop
14
                Put (H'Image & " ");
15
              end loop;
16
             Put_Line (M'Last (2)'Image);
17
18
          end loop;
       end Show_Indices;
19
20
       procedure Show_Values (M : Measurements) is
21
          package H_IO is
22
            new Ada.Text_I0.Integer_I0 (Hours);
23
          package M IO is
24
            new Ada.Text_IO.Float_IO (Measurement);
25
26
            procedure Set_I0_Defaults is
27
28
            begin
                H_IO.Default_Width := 5;
29
30
               M_IO.Default_Fore := 1;
31
               M_IO.Default_Aft
                                     := 2;
32
               M_IO.Default_Exp
                                     := 0;
33
            end Set_I0_Defaults;
34
       begin
35
          Set_IO_Defaults;
36
37
          Put_Line ("---- Values ----");
38
          Put (" ");
39
          for H in M'Range (2) loop
40
             H_IO.Put (H);
41
          end loop;
42
          New_Line;
43
44
          for D in M'Range (1) loop
45
             Put (D'Image & " ");
46
47
              for H in M'Range (2) loop
48
                M_IO.Put (M (D, H));
49
                 Put (" ");
50
             end loop;
51
             New_Line;
52
          end loop;
53
       end Show_Values;
54
55
       procedure Reset (M : out Measurements) is
56
       begin
57
          M := (others => (others => 0.0));
58
       end Reset;
59
60
   end Measurement_Defs;
61
```

4

Listing 11:	show	measurements.adb

```
with Measurement_Defs; use Measurement_Defs;
1
2
  procedure Show Measurements is
3
      M : Measurements;
4
  begin
5
      Reset (M);
6
      Show_Indices (M);
7
      Show_Values (M);
8
  end Show Measurements;
g
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Multidimensional_Arrays.

→Multidimensional_Measurements

MD5: bcffa3913007bd9152149ad9616842b8
```

Runtime output

	· Ind	ic	es ·												
MON	0	1	2	3	4	5	67	7	89	10	11				
TUE	0	1	2	3	4	5	67	7	89	10	11				
WED	0	1	2	3	4	5	67	7	89	10	11				
THU	0	1	2	3	4	5	67	7	89	10	11				
FRI	0	1	2	3	4	5	67	7	89	10	11				
SAT	0	1	2	3	4	5	67	7	89	10	11				
SUN	0	1	2	3	4	5	67	7	89	10	11				
	- Val	ue	s		•										
	C		1		2	3		4	5	6	7	8	9	10	11
MON	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00
TUE	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00
WED	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00
THU	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00
FRI	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SAT	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SUN	0.00	0	.00	0.	00	0.00	0.0)0	0.00	0.00	0.00	0.00	0.00	0.00	0.00

We recommend that you spend some time analyzing this example. Also, we'd like to highlight the following aspects:

- We access a value from a multidimensional array by using commas to separate the index values within the parentheses. For example: M (D, H) allows us to access the value on day D and hour H from the multidimensional array M.
- To loop over the multidimensional array M, we write **for** D **in** M'Range (1) **loop** and **for** H **in** M'Range (2) **loop** for the first and second dimensions, respectively.
- To reset all values of the multidimensional array, we use an aggregate with this form: (others => (others => 0.0)).

In the Ada Reference Manual

- 3.6 Array Types¹²⁰
- 3.6.2 Operations of Array Types¹²¹

¹²⁰ http://www.ada-auth.org/standards/22rm/html/RM-3-6.html

121 http://www.ada-auth.org/standards/22rm/html/RM-3-6-2.html

6.2.1 Unconstrained Multidimensional Arrays

Previously, we've discussed unconstrained arrays for the unidimensional case. It's possible to declare unconstrained multidimensional arrays as well. For example:

```
Listing 12: multidimensional_arrays_decl.ads
```

```
package Multidimensional Arrays Decl is
1
2
      type F1 is array (Positive range <>) of Float;
3
      type F2 is array (Positive range <>,
4
                          Positive range <>) of Float;
5
      type F3 is array (Positive range <>,
6
                          Positive range <>,
7
                         Positive range <>) of Float;
8
9
   end Multidimensional_Arrays_Decl;
10
```

Code block metadata

Here, we're declaring the one-dimensional type F1, the two-dimensional type F2 and the three-dimensional type F3.

As is the case with unidimensional arrays, we must specify the bounds when declaring objects of unconstrained multidimensional array types:

Listing 13: show_multidimensional_arrays.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Multidimensional Arrays Decl;
3
   use Multidimensional Arrays Decl;
4
5
   procedure Show Multidimensional Arrays is
6
      A1 : F1 (1 ... 2);
7
      A2 : F2 (1 ... 4, 10 ... 20);
8
      A3 : F3 (2 .. 3, 1 .. 5, 1 .. 2);
9
   begin
10
      Put Line ("A1'Length (1): "
11
                 & A1'Length (1)'Image);
12
      Put_Line ("A2'Length (1): "
13
                 & A2'Length (1)'Image);
14
      Put_Line ("A2'Length (2): "
15
                 & A2'Length (2)'Image);
16
      Put_Line ("A3'Length (1):
17
                 & A3'Length (1)'Image);
18
      Put_Line ("A3'Length (2): "
19
                 & A3'Length (2)'Image);
20
      Put_Line ("A3'Length (3): "
21
                 & A3'Length (3)'Image);
22
   end Show_Multidimensional_Arrays;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Multidimensional_Arrays.

GUnconstrained_Multidimensional_Arrays

MD5: 9fb007abbfe238345d80cb315bb834c9
```

Build output

Runtime output

Al'Length (1): 2 A2'Length (1): 4 A2'Length (2): 11 A3'Length (1): 2 A3'Length (2): 5 A3'Length (3): 2

6.2.2 Arrays of arrays

It's important to distinguish between multidimensional arrays and arrays of arrays. Both are supported in Ada, but they're very distinct from each other. We can create an array of an array by first specifying a one-dimensional array type T1, and then specifying another one-dimensional array type T2 where each component of T2 is of T1 type:

Listing 14: array_of_arrays_decl.ads

```
package Array_Of_Arrays_Decl is
1
2
      type T1 is
3
        array (Positive range <>) of Float;
4
5
      type T2 is
6
        array (Positive range <>) of T1 (1 .. 10);
7
8
                                    bounds must be set!
9
10
   end Array_Of_Arrays_Decl;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Arrays.Array_Of_Arrays.Array_Of_Arrays
MD5: fd67739bb21f202615180aa02f5284aa

Note that, in the declaration of T2, we must set the bounds for the T1 type. This is a major difference to multidimensional arrays, which allow for unconstrained ranges in multiple dimensions.

We can rewrite the previous application for measurements using arrays of arrays. This is the adapted code:

Listing 15: measurement defs.ads

```
package Measurement Defs is
1
2
      type Days is
3
         (Mon, Tue, Wed, Thu, Fri, Sat, Sun);
4
5
      type Hours is range 0 .. 11;
6
7
      subtype Measurement is Float;
8
9
      type Hourly Measurements is
10
```

```
array (Hours) of Measurement;
11
12
       type Measurements is
13
         array (Days) of Hourly_Measurements;
14
15
       procedure Show_Indices (M : Measurements);
16
17
       procedure Show_Values (M : Measurements);
18
19
       procedure Reset (M : out Measurements);
20
21
   end Measurement_Defs;
22
```

Listing 16: measurement_defs.adb

```
with Ada.Text_I0;
                             use Ada.Text_I0;
1
2
   package body Measurement Defs is
3
4
       procedure Show_Indices (M : Measurements) is
5
       begin
6
          Put_Line ("---- Indices ----");
7
8
          for D in M'Range loop
9
             Put (D'Image & " ");
10
11
             for H in M (D) 'First ...
12
                       M (D) 'Last - 1
13
             loop
14
                 Put (H'Image & " ");
15
             end loop;
16
             Put_Line (M (D)'Last'Image);
17
          end loop;
18
       end Show_Indices;
19
20
       procedure Show Values (M : Measurements) is
21
          package H IO is
22
            new Ada.Text IO.Integer IO (Hours);
23
          package M IO is
24
            new Ada.Text_IO.Float_IO (Measurement);
25
26
            procedure Set_IO_Defaults is
27
28
            begin
               H_IO.Default_Width := 5;
29
30
               M_IO.Default_Fore := 1;
31
               M_IO.Default_Aft
                                     := 2;
32
               M IO.Default Exp
                                     := 0;
33
            end Set_I0_Defaults;
34
       begin
35
          Set_I0_Defaults;
36
37
          Put_Line ("---- Values ----");
38
          Put (" ");
39
          for H in M (M'First)'Range loop
40
             H_IO.Put (H);
41
          end loop;
42
          New Line;
43
44
          for D in M'Range loop
45
             Put (D'Image & " ");
46
```

```
47
              for H in M (D)'Range loop
48
                 M_IO.Put (M (D) (H));
49
                 Put (" ");
50
             end loop;
51
             New_Line;
52
          end loop;
53
       end Show_Values;
54
55
       procedure Reset (M : out Measurements) is
56
       begin
57
          M := (others => (others => 0.0));
58
59
       end Reset;
60
   end Measurement_Defs;
61
```

Listing 17: show_measurements.adb

```
with Measurement_Defs; use Measurement_Defs;
1
2
   procedure Show_Measurements is
3
      M : Measurements;
4
  begin
5
      Reset (M);
6
      Show Indices (M);
7
      Show Values (M);
8
  end Show Measurements;
9
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Arrays.Array_Of_Arrays.Multidimensional_ →Measurements MD5: 5cb66bbb1890787b7c023406b2cafb4d

Runtime output

```
---- Indices ----
MON 0 1 2
    34
       5
         7
           9
        6
          8
            10
              11
TUE 0 1 2
    3
       5
         7
           9
     4
        6
          8
            10
              11
         7
WED 0 1 2
    34
       5
          8
           9
        6
            10
             11
       5
         7
THU 0 1 2
    34
        6
          8
           9
            10
             11
FRI 0 1 2
    34
       5
        6
         7
          8
           9
            10
             11
SAT 0
  1
   2
    3
      4
       5
        6
         7
          8
           9
            10
              11
SUN
 0
  1
   23
      4
       5
        6
         7
          8
           9
            10
             11
 Values -
           5
  0
    1
      2
       3
         4
             6
               7
                8
                  g
                    10
                     11
```

Again, we recommend that you spend some time analyzing this example and comparing it to the previous version that uses multidimensional arrays. Also, we'd like to highlight the following aspects:

We access a value from an array of arrays by specifying the index of each array separately. For example: M (D) (H) allows us to access the value on day D and hour H from the array of arrays M.

- To loop over an array of arrays M, we write **for** D **in** M'Range **loop** for the first level of M and **for** H **in** M (D)'Range **loop** for the second level of M.
- Resetting all values of an array of arrays is very similar to how we do it for multidimensional arrays. In fact, we can still use an aggregate with this form: (others => (others => 0.0)).

6.3 Derived array types and array subtypes

6.3.1 Derived array types

As expected, we can derive from array types by declaring a new type. Let's see a couple of examples based on the Measurement_Defs package from previous sections:

```
Listing 18: measurement_defs.ads
```

```
package Measurement Defs is
1
2
      type Measurements is
3
        array (Positive range <>) of Float;
4
5
6
       -- New array type:
7
8
9
      type Measurements Derived is
10
        new Measurements;
11
12
      -- New array type with
13
      -- default component value:
14
15
      type Measurements Def30 is
16
        new Measurements
17
          with Default Component Value => 30.0;
18
19
20
      -- New array type with constraints:
21
22
      type Measurements_10 is
23
        new Measurements (1 .. 10);
24
25
   end Measurement Defs;
26
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Derived_Arrays_And_Subtypes.

→Derived_Arrays

MD5: aefef9b9a844ad820d7f16546b8ffa64
```

In this example, we're deriving Measurements_Derived from the Measurements type. In the case of the Measurements_Def30 type, we're not only deriving from the Measurements type, but also setting the *default component value* (page 66) to 30.0. Finally, in the case of the Measurements_10, we're deriving from the Measurements type and *constraining the array type* (page 296) in the range from 1 to 10.

Let's use these types in a test application:

Listing 19: show_measurements.adb

```
with Measurement_Defs; use Measurement_Defs;
```

2

```
procedure Show_Measurements is
3
      M1, M2 : Measurements (1 .. 10)
4
                   := (others => 0.0);
5
6
            : Measurements_Derived (1 .. 10);
      MD
7
             : Measurements_Derived (1 .. 40);
      MD2
8
      MD10 : Measurements 10;
9
   begin
10
      Μ1
           := M2;
11
           ~~~~~
       - -
12
       -- Assignment of arrays of
13
      -- same type.
14
15
      MD := Measurements_Derived (M1);
16
               ~~~~~~~~~~~<u>~</u>~~~~~
17
      - -
       -- Conversion to derived type for
18
      -- the assignment.
19
20
      MD10 := Measurements 10 (M1);
21
               ~~~~~~
22
       -- Conversion to derived type for
23
       -- the assignment.
24
25
      MD10 := Measurements_10 (MD);
26
      MD10 := Measurements_10 (MD2 (1 .. 10));
27
   end Show_Measurements;
28
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Arrays.Derived_Arrays_And_Subtypes.

→Derived_Arrays

MD5: ce37a9c17eb9e1bb3931cca82852b54a
```

Build output

```
show_measurements.adb:8:04: warning: variable "MD2" is read but never assigned [-
_gnatwv]
```

As illustrated by this example, we can assign objects of different array types, provided that we perform the appropriate type conversions and make sure that the bounds match.

6.3.2 Array subtypes

Naturally, we can also declare subtypes of array types. For example:

```
Listing 20: measurement_defs.ads
```

```
package Measurement_Defs is
1
2
      type Measurements is
3
        array (Positive range <>) of Float;
4
5
6
       -- Simple subtype declaration:
7
8
      subtype Measurements Sub is Measurements;
9
10
11
      -- Subtype with constraints:
12
13
```

```
subtype Measurements_10 is
14
         Measurements (1 .. 10);
15
16
17
          Subtype with dynamic predicate
       - -
18
           (array can only have 20 components
       - -
19
           at most):
       - -
20
21
       subtype Measurements Max 20 is Measurements
22
           with Dynamic Predicate =>
23
                   Measurements Max 20'Length <= 20;
24
25
26
          Subtype with constraints and
27
       - -
           dynamic predicate (first element
28
       - -
       -- must be 2.0).
29
30
      subtype Measurements_First_Two is
31
         Measurements (1 .. 10)
32
           with Dynamic Predicate =>
33
                   Measurements_First_Two (1) = 2.0;
34
35
   end Measurement_Defs;
36
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Arrays.Derived_Arrays_And_Subtypes.Array_ →Subtypes MD5: fa03c836111aa2df223a38a5d04d18bc

Here, we're declaring subtypes of the Measurements type. For example, Measurements_Sub is a *simple* subtype of Measurements type. In the case of the Measurements_10 subtype, we're constraining the type to a range from 1 to 10.

For the Measurements_Max_20 subtype, we're specifying — via a dynamic predicate — that arrays of this subtype can only have 20 components at most. Finally, for the Measurements_First_Two subtype, we're constraining the type to a range from 1 to 10 and requiring that the first component must have a value of 2.0.

Note that we cannot set the default component value for array subtypes — only type declarations are allowed to use that facility.

Let's use these subtypes in a test application:

```
Listing 21: show_measurements.adb
```

```
with Measurement_Defs; use Measurement_Defs;
1
2
   procedure Show Measurements is
3
      M1, M2 : Measurements (1 .. 10)
4
                   := (others => 0.0);
5
      MS
               : Measurements_Sub (1 .. 10);
6
               : Measurements_10;
      MD10
7
      M_Max20 : Measurements_Max_20 (1 .. 40);
8
               : Measurements_First_Two;
      M F2
9
   begin
10
      MS
                := M1;
11
      MD10
                := M1;
12
13
      M_Max20 := (others => 0.0); -- ERROR!
14
15
      MD10(1) := 4.0;
16
```

17	M_F2	:= MD10;	 ERROR!
18	end Show	Measurements;	

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Arrays.Derived_Arrays_And_Subtypes.Array_ ⇔Subtypes MD5: 003ddaab65d8c163302811abd7889745

Runtime output

raised ADA.ASSERTIONS.ASSERTION_ERROR : Dynamic_Predicate failed at show_ →measurements.adb:14

As expected, assignments to objects with different subtypes — but with the same parent type — work fine without conversion. The assignment to M_Max_20 fails because of the predicate failure: the predicate requires that the length be 20 at most, and it's 40 in this case. Also, the assignment to M_F2 fails because the predicate requires that the first element must be set to 2.0, and MD10 (1) has the value 4.0.

CHAPTER SEVEN

STRINGS

7.1 Character and String Literals

So far, we're already seen many examples of string literals — both in the Introduction to Ada^{122} course and in the present course. In this section, we define them once more and discuss a couple of details about them.

7.1.1 Character Literals

A character literal is simply a character between apostrophes (or *single quotation marks*). For example:

Listing 1: show charac	ter literals.adb
------------------------	------------------

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
  procedure Show_Character_Literals is
3
      C : Character := 'a';
4
5
      - -
                   Character literal
6
  begin
7
      Put Line ("Character : " & C);
8
  end Show Character Literals;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Character_String_Literals.

⇔Character_Literals

MD5: e9bf0dee97b4c6d52937316e7f285f48
```

Runtime output

Character : a

1

In this example, we initialize the character variable C with the character literal 'a'.

7.1.2 String Literals

A string literal is simply a collection of characters between quotation marks. For example:

Listing 2: show simple string literals.adb

```
with Ada.Text_I0; use Ada.Text_I0;
```

```
procedure Show_Simple_String_Literals is
```

(continues on next page)

 $^{122}\ https://learn.adacore.com/courses/intro-to-ada/index.html \# intro-ada-course-index$

```
S1 : String := "Hello";
4
                      ~~~~~
5
      - -
      - -
                 String literal
6
7
      S2 : String := "World";
8
                      ~~~~~
9
      - -
      - -
                   String literal
10
  begin
11
      Put_Line (S1 & " " & S2);
12
  end Show_Simple_String_Literals;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Character_String_Literals.Simple_

→String_Literals

MD5: a19bfa1ab6048f8cad9858d57b9f21e1
```

Runtime output

Hello World

In this example, "Hello" and "World" are string literals.

String literals with quotation

If you want to include a quotation mark in a string literal, you have to write "" (inside that string literal):

Listing 3: show_string_literals_with_quotes.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_String_Literals_With_Quotes is
3
      S1 : String := "Hello";
4
      S2 : String := "World";
5
   begin
6
      7
                   ~~
8
      - -
              Quotation marks
9
      - -
                & " " & S2 & " "" ");
10
                                 \wedge \wedge
11
       - -
      - -
               Quotation marks
12
13
      Put_Line ("""Hello World!""");
14
                  ~~
15
       - -
                  Quotation marks
16
17
      Put_Line ("""""");
18
                  ~~~~
19
      - -
       - -
            Quotation marks
20
  end Show_String_Literals_With_Quotes;
21
```

Code block metadata

Runtime output

" Hello World " "Hello World!" ""

In this example, we display "Hello World "to the user by adding quotation marks to the concatenated strings in the call to Put_Line.

Note that the three quotation marks at the beginning of """Hello World!""" consist of the quotation mark that indicate the beginning of the string literal and the two quotation marks that represent a single quotation mark inside the string literal. (The same thing happens at the end of this string literal, but in reverse.) This string literal is displayed as "Hello World!" to the user.

Finally, the string literal """"" is displayed as "" to the user.

Empty string literals

An empty string is represented by quotation marks without characters in between: "". For example:

Listing 4: show_empty_string_literals.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Empty String Literals is
3
      S1 : String
4
                           :=
      S2 : String (1 .. 0) := "";
5
   begin
6
      Put_Line (S1);
7
      Put_Line (S2);
8
      Put_Line ("");
9
  end Show_Empty_String_Literals;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Character_String_Literals.Empty_

→String_Literals

MD5: f2f7c47784f1053665db9499cb6b53d8
```

Runtime output

Note that an empty string is an array of characters without any components. This is made explicit by the declaration of S2. Here, by using the range $1 \ldots 0$, we're declaring an empty array.

In other languages

In C, an empty string still contains a single character: the null character ($\0$). In Ada, however, an empty string doesn't have any characters.

In the Ada Reference Manual

- 2.5 Character Literals¹²³
- 2.6 String Literals¹²⁴

7.2 Wide and Wide-Wide Strings

We've seen many source-code examples so far that includes strings. In most of them, we were using the standard string type: **String**. This type is useful for the common use-case of displaying messages or dealing with information in plain English. Here, we define "plain English" as the use of the language that avoids French accents or German umlaut, for example, and doesn't make use of any characters in non-Latin alphabets.

There are two additional string types in Ada: Wide_String, and Wide_Wide_String. These types are particularly important when dealing with textual information in non-standard English, or in various other languages, non-Latin alphabets and special symbols.

These string types use different bit widths for their characters. This becomes more apparent when looking at the type definitions:

```
type String is
array (Positive range <>) of Character;
type Wide_String is
array (Positive range <>) of Wide_Character;
type Wide_Wide_String is
array (Positive range <>) of
Wide_Wide_Character;
```

The following table shows the typical bit-width of each character of the string types:

Character Type	Width	
Character	8 bits	
Wide_Character	16 bits	
Wide Wide Character	32 bits	

We can see that when running this example:

Listing 5: show_wide_char_types.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
3
   procedure Show_Wide_Char_Types is
4
   begin
      Put Line ("Character'Size:
5
                 & Integer'Image
6
                      (Character'Size));
7
      Put_Line ("Wide_Character'Size:
                                              n.
8
                 & Integer'Image
9
                      (Wide Character'Size));
10
      Put_Line ("Wide_Wide_Character'Size:
11
                 & Integer'Image
12
                      (Wide_Wide_Character'Size));
13
   end Show_Wide_Char_Types;
14
```

Code block metadata

123 http://www.ada-auth.org/standards/22rm/html/RM-2-5.html

124 http://www.ada-auth.org/standards/22rm/html/RM-2-6.html

Runtime output

Character'Size:	8
Wide_Character'Size:	16
Wide_Wide_Character'Size:	32

Let's look at another example, this time using wide strings:

Listing 6: show_wide_string_types.adb

```
with Ada.Text_I0;
1
   with Ada.Wide_Text_I0;
2
   with Ada.Wide_Wide_Text_I0;
3
4
   procedure Show_Wide_String_Types is
5
      package TI renames Ada.Text_I0;
package WTI renames Ada.Wide_Text_I0;
6
7
      package WWTI renames Ada.Wide Wide Text I0;
8
9
                                        := "hello";
      S
          : constant String
10
      WS : constant Wide String := "hello";
11
      WWS : constant Wide_Wide_String := "hello";
12
13
   begin
                                       "&S);
      TI.Put_Line ("String:
14
      TI.Put_Line ("Length:
                                        .....
15
                    & Integer'Image (S'Length));
16
      TI.Put_Line ("Size:
17
                    & Integer'Image (S'Size));
18
      TI.Put Line ("Component Size:
19
                    & Integer'Image
20
                       (S'Component_Size));
21
      TI.Put_Line ("-----");
22
23
      WTI.Put_Line ("Wide string:
                                       " & WS);
24
                                       .....
      TI.Put_Line ("Length:
25
                   & Integer'Image (WS'Length));
26
      TI.Put_Line ("Size:
27
                    & Integer'Image (WS'Size));
28
      TI.Put_Line ("Component_Size:
29
                   & Integer'Image
30
                       (WS'Component_Size));
31
      TI.Put_Line ("-----");
32
33
      WWTI.Put_Line ("Wide-wide string: " & WWS);
34
      TI.Put_Line ("Length:
35
                    & Integer'Image (WWS'Length));
36
      TI.Put_Line ("Size:
37
                    & Integer'Image (WWS'Size));
38
      TI.Put_Line ("Component_Size:
39
                   & Integer'Image
40
                        (WWS'Component_Size));
41
      TI.Put Line ("-----");
42
   end Show_Wide_String_Types;
43
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Wide_Wide-Wide_Strings.Wide_ →String_Types MD5: 137816c6fd78add34287a72e45cf4fb7

Runtime output

String: hello 5 Length: 40 Size: Component_Size: 8 - - -Wide string: hello Lenath: 5 Size: 80 Component Size: 16 Wide-wide string: hello Length: 5 Size: 160 Component Size: 32

Here, all strings (S, WS and WWS) have the same length of 5 characters. However, the size of each character is different — thus, each string has a different overall size.

The recommendation is to use the **String** type when the textual information you're processing is in standard English. In case any kind of internationalization is needed, using Wide_Wide_String is probably the best choice, as it covers all possible use-cases.

1 In the Ada Reference Manual

• 3.6.3 String Types¹²⁵

7.2.1 Text I/O

Note that, in the previous example, we were using different versions of the Ada.Text_I0 package depending on the string type we were using:

- Ada.Text_IO for objects of String type,
- Ada.Wide_Text_IO for objects of Wide_String type,
- Ada.Wide_Wide_Text_IO for objects of Wide_Wide_String type.

In that example, we were also using package renaming to differentiate among those packages.

Similarly, there are different versions of text I/O packages for individual types. For example, if we want to display the value of a Long_Integer variable based on the Wide_Wide_String type, we can select the Ada.Long_Integer_Wide_Wide_Text_I0 package. In fact, the list of packages resulting from the combination of those types is quite long:

¹²⁵ http://www.ada-auth.org/standards/22rm/html/RM-3-6-3.html

Scalar Type	Text I/O Packages
Integer	 Ada.Integer_Text_I0 Ada.Integer_Wide_Text_I0 Ada.Integer_Wide_Wide_Text_I0
Long_Integer	 Ada.Long_Integer_Text_I0 Ada.Long_Integer_Wide_Text_I0 Ada.Long_Integer_Wide_Wide_Text_I0
Long_Long_Integer	 Ada.Long_Long_Integer_Text_I0 Ada.Long_Long_Integer_Wide_Text_I0 Ada.Long_Long_Integer_Wide_Wide_Text_I0
Float	 Ada.Float_Text_I0 Ada.Float_Wide_Text_I0 Ada.Float_Wide_Wide_Text_I0
Long_Float	 Ada.Long_Float_Text_I0 Ada.Long_Float_Wide_Text_I0 Ada.Long_Float_Wide_Wide_Text_I0
Long_Long_Float	 Ada.Long_Long_Float_Text_I0 Ada.Long_Long_Float_Wide_Text_I0 Ada.Long_Long_Float_Wide_Wide_Text_I0

Also, there are different versions of the generic packages Integer_I0 and Float_I0:

Scalar Type	Text I/O Packages
Integer types	 Ada.Text_I0.Integer_I0 Ada.Wide_Text_I0.Integer_I0 Ada.Wide_Wide_Text_I0. Integer_I0
Real types	 Ada.Text_I0.Float_I0 Ada.Wide_Text_I0.Float_I0 Ada.Wide_Wide_Text_I0.Float_I0

1 In the Ada Reference Manual

- A.10 Text Input-Output¹²⁶
- A.10.1 The Package Text_IO¹²⁷
- A.10.8 Input-Output for Integer Types¹²⁸

- A.10.9 Input-Output for Real Types¹²⁹
- A.11 Wide Text Input-Output and Wide Wide Text Input-Output¹³⁰

7.2.2 Wide and Wide-Wide String Handling

As we've just seen, we have different versions of the Ada.Text_IO package. The same applies to string handling packages. As we've seen in the Introduction to Ada course¹³¹, we can use the Ada.Strings.Fixed and Ada.Strings.Maps packages for string handling. For other formats, we have these packages:

- Ada.Strings.Wide_Fixed,
- Ada.Strings.Wide_Wide_Fixed,
- Ada.Strings.Wide_Maps,
- Ada.Strings.Wide_Wide_Maps.

Let's look at this example¹³² from the Introduction to Ada course, which we adapted for wide-wide strings:

Listing	7:	show	find	words.adb
---------	----	------	------	-----------

```
with Ada.Strings; use Ada.Strings;
1
2
   with Ada.Strings.Wide_Wide_Fixed;
3
   use Ada.Strings.Wide_Wide_Fixed;
4
5
   with Ada.Strings.Wide_Wide_Maps;
6
   use Ada.Strings.Wide_Wide_Maps;
7
8
   with Ada.Wide Wide Text IO;
9
   use Ada.Wide Wide Text IO;
10
11
   procedure Show_Find_Words is
12
13
           : constant Wide_Wide_String :=
       S
14
               "Hello" & 3 * " World";
15
      F
         : Positive;
16
          : Natural;
17
      1
          : Natural := 1;
      Ι
18
19
      Whitespace : constant
20
         Wide Wide Character Set :=
21
           To_Set (' ');
22
   begin
23
      Put_Line ("String: " & S);
24
      Put_Line ("String length: "
25
                 & Integer'Wide Wide Image
26
                      (S'Length));
27
28
      while I in S'Range loop
29
```

(continues on next page)

¹²⁶ http://www.ada-auth.org/standards/22rm/html/RM-A-10.html

¹²⁷ http://www.ada-auth.org/standards/22rm/html/RM-A-10-1.html

¹²⁸ http://www.ada-auth.org/standards/22rm/html/RM-A-10-8.html

¹²⁹ http://www.ada-auth.org/standards/22rm/html/RM-A-10-9.html

¹³⁰ http://www.ada-auth.org/standards/22rm/html/RM-A-11.html

 $^{^{131}}$ https://learn.adacore.com/courses/intro-to-ada/chapters/standard_library_strings.html# intro-ada-string-operations

¹³² https://learn.adacore.com/courses/intro-to-ada/chapters/standard_library_strings.html# intro-ada-string-operations-show-find-words

```
Find_Token
30
             (Source => S,
31
                      => Whitespace,
              Set
32
             From
                      => I,
33
             Test
                      => Outside,
34
             First
                      => F,
35
             Last
                      => L);
36
37
          exit when L = 0;
38
39
          Put Line ("Found word instance at position "
40
                     & F'Wide_Wide_Image
41
                     & ": '" & S (F .. L) & "'");
42
43
          I := L + 1;
44
       end loop;
45
46
   end Show_Find_Words;
47
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Wide_Wide_Wide_Strings.Wide_Wide_ ⇔String_Handling MD5: 3b5a4d61e6dc5bd16e85f85580ad82ae

Runtime output

```
String: Hello World World World
String length: 23
Found word instance at position 1: 'Hello'
Found word instance at position 7: 'World'
Found word instance at position 13: 'World'
Found word instance at position 19: 'World'
```

In this example, we're using the Find_Token procedure to find the words from the phrase stored in the S constant. All the operations we're using here are similar to the ones for **String** type, but making use of the Wide_Wide_String type instead. (We talk about the Wide_Wide_Image attribute *later on* (page 339).)

In the Ada Reference Manual

- A.4.6 String-Handling Sets and Mappings¹³³
- A.4.7 Wide_String Handling¹³⁴
- A.4.8 Wide_Wide_String Handling¹³⁵

7.2.3 Bounded and Unbounded Wide and Wide-Wide Strings

We've seen in the Introduction to Ada course that other kinds of **String** types are available. For example, we can use bounded¹³⁶ and unbounded strings¹³⁷ — those correspond to the Bounded_String and Unbounded_String types.

¹³³ http://www.ada-auth.org/standards/22rm/html/RM-A-4-6.html

¹³⁴ http://www.ada-auth.org/standards/22rm/html/RM-A-4-7.html

¹³⁵ http://www.ada-auth.org/standards/22rm/html/RM-A-4-8.html

 $^{^{136}\} https://learn.adacore.com/courses/intro-to-ada/chapters/standard_library_strings.html# intro-ada-bounded-strings$

¹³⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/standard_library_strings.html# intro-ada-unbounded-strings

Those kinds of string types are available for **Wide_String**, and Wide_Wide_String. The following table shows the available types and corresponding packages:

Туре	Package
Bounded Wide String	Ada.Strings.Wide Bounded
Bounded_Wide_Wide_String	Ada.Strings.Wide_Wide_Bounded
Unbounded_Wide_String	Ada.Strings.Wide_Unbounded
Unbounded_Wide_Wide_String	Ada.Strings.Wide_Wide_Unbounded

The same applies to text I/O for those strings. For the standard case, we have Ada. Text_I0.Bounded_I0 for the Bounded_String type and Ada.Text_I0.Unbounded_I0 for the Unbounded_String type.

For wider string types, we have:

Туре	Text I/O Package
Bounded_Wide_String Bounded Wide Wide String	Ada.Wide_Text_IO.Wide_Bounded_IO Ada.Wide Wide Text IO.Wide Wide Bounded IO
Unbounded_Wide_String	Ada.Wide_Text_I0.Wide_Unbounded_I0
Unbounded_Wide_Wide_String	Ada.Wide_Wide_Text_IO.Wide_Wide_Unbounded_IO

Let's look at a simple example:

Listing 8: show_unbounded_wide_wide_string.adb

```
with Ada.Strings.Wide Wide Unbounded;
1
   use Ada.Strings.Wide_Wide_Unbounded;
2
3
   with Ada.Wide_Wide_Text_IO.Wide_Wide_Unbounded_IO;
4
   use Ada.Wide_Wide_Text_IO.Wide_Wide_Unbounded_IO;
5
6
   procedure Show Unbounded Wide Wide String is
7
      S : Unbounded Wide Wide String
8
        := To Unbounded Wide Wide String ("Hello");
9
   begin
10
      S := S & Wide_Wide_String'(" hello");
11
      Put_Line ("Unbounded wide-wide string: " & S);
12
   end Show_Unbounded_Wide_Wide_String;
13
```

Code block metadata

Runtime output

Unbounded wide-wide string: Hello hello

In this example, we're declaring a variable S and initializing it with the word "Hello." Then, we're concatenating it with " hello" and displaying it. All the operations we're using here are similar to the ones for Unbounded_String type, but they've been adapted for the Unbounded_Wide_Wide_String type.

In the Ada Reference Manual

- A.4.7 Wide_String Handling¹³⁸
- A.4.8 Wide_Wide_String Handling¹³⁹
- A.11 Wide Text Input-Output and Wide Wide Text Input-Output¹⁴⁰

7.3 String Encoding

Unicode is one of the most widespread standards for encoding writing systems other than the Latin alphabet. It defines a format called Unicode Transformation Format (UTF)¹⁴¹ in various versions, which vary according to the underlying precision, support for backwards-compatibility and other requirements.

1 In the Ada Reference Manual

• A.4.11 String Encoding¹⁴²

7.3.1 UTF-8 encoding and decoding

A common UTF format is UTF-8, which encodes strings using up to four (8-bit) bytes and is backwards-compatible with the ASCII format. While encoding of ASCII characters requires only one byte, Chinese characters require three bytes, for example.

In Ada applications, UTF-8 strings are indicated by using the UTF_8_String from the Ada. Strings.UTF_Encoding package. In order to encode from and to UTF-8 strings, we can use the Encode and Decode functions. Those functions are specified in the child packages of the *Ada.Strings.UTF_Encoding* package. We select the appropriate child package depending on the string type we're using, as you can see in the following table:

Child Package of Ada.Strings.UTF_Encoding	Convert from / to
.Strings	String type
.Wide_Strings	Wide_String type
.Wide_Wide_Strings	Wide_Wide_String type

Let's look at an example:

Listing 9: show ww utf string.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Ada.Strings.UTF_Encoding;
Ada.Strings.UTF_Encoding;
with Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
ause Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
with Ada.Strings.Wide_Wide_Unbounded;
```

(continues on next page)

¹³⁸ http://www.ada-auth.org/standards/22rm/html/RM-A-4-7.html

¹⁴¹ https://unicode.org/faq/utf_bom.html#gen2

¹³⁹ http://www.ada-auth.org/standards/22rm/html/RM-A-4-8.html

¹⁴⁰ http://www.ada-auth.org/standards/22rm/html/RM-A-11.html

¹⁴² http://www.ada-auth.org/standards/22rm/html/RM-A-4-11.html

use Ada.Strings.Wide_Wide_Unbounded;

10 11 (continued from previous page)

```
procedure Show_WW_UTF_String is
12
13
      function To_UWWS
14
         (Source : Wide_Wide_String)
15
          return Unbounded_Wide_Wide_String
16
            renames To_Unbounded_Wide_Wide_String;
17
18
      function To_WWS
19
         (Source : Unbounded Wide Wide String)
20
          return Wide_Wide_String
21
22
            renames To_Wide_Wide_String;
23
      Hello_World_Arabic : constant
24
        :عالم" يا مرحبا" =: UTF_8_String;
25
      WWS_Hello_World_Arabic : constant
26
        Wide_Wide_String :=
27
           Decode (Hello_World_Arabic);
28
29
      UWWS : Unbounded_Wide_Wide_String;
30
   begin
31
      UWWS := "Hello World: "
32
               & To_UWWS (WWS_Hello_World_Arabic);
33
34
      Show_WW_String : declare
35
         WWS : constant Wide_Wide_String :=
36
                  To_WWS (UWWS);
37
      begin
38
          Put_Line ("Wide_Wide_String Length: "
39
                    & WWS'Length'Image);
40
          Put Line ("Wide Wide String Size:
41
                    & WWS'Size'Image);
42
      end Show_WW_String;
43
44
      Put_Line
45
                          -----");
         ("----
46
      Put_Line
47
         ("Converting Wide_Wide_String to UTF-8...");
48
49
      Show_UTF_8_String : declare
50
          S UTF 8 : constant UTF 8 String :=
51
                      Encode (To_WWS (UWWS));
52
      begin
53
         Put Line ("UTF-8 String:
54
                    & S_UTF_8);
55
          Put_Line ("UTF-8 String Length: "
56
                    & S_UTF_8'Length'Image);
57
          Put_Line ("UTF-8 String Size:
58
                    & S UTF 8'Size'Image);
59
      end Show_UTF_8_String;
60
61
   end Show WW UTF String;
62
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.WW_UTF_String
MD5: cecfb420bb804f42e7a65b793abcbef5
```

Runtime output

In this application, we start by storing a string in Arabic in the Hello_World_Arabic constant. We then use the Decode function to convert that string from UTF_8_String type to Wide_Wide_String type — we store it in the WWS_Hello_World_Arabic constant.

We use a variable of type Unbounded_Wide_Wide_String (UWWS) to manipulate strings: we append the string in Arabic to the "Hello World: " string and store it in UWWS.

In the Show_WW_String block, we convert the string — stored in UWWS — from the Unbounded_Wide_Wide_String type to the Wide_Wide_String type and display the length and size of the string. We do something similar in the Show_UTF_8_String block, but there, we convert to the UTF_8_String type.

Also, in the Show_UTF_8_String block, we use the Encode function to convert that string from Wide_Wide_String type to then UTF_8_String type — we store it in the S_UTF_8 constant.

7.3.2 UTF-8 size and length

As you can see when running the last code example from the previous subsection, we have different sizes and lengths depending on the string type:

String type	Size	Length
Wide_Wide_String	832	26
UTF_8_String	296	37

The size needed for storing the string when using the Wide_Wide_String type is bigger than the one when using the UTF_8_String type. This is expected, as the Wide_Wide_String uses 32-bit characters, while the UTF_8_String type uses 8-bit codes to store the string in a more efficient way (memory-wise).

The length of the string using the Wide_Wide_String type is equivalent to the number of symbols we have in the original string: 26 characters / symbols. When using UTF-8, however, we may need more 8-bit codes to represent one symbol from the original string, so we may end up with a length value that is bigger than the actual number of symbols from the original string — as it is the case in this source-code example.

This difference in sizes might not always be the case. In fact, the sizes match when encoding a symbol in UTF-8 that requires four 8-bit codes. For example:

Listing 10: show_utf_8.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Strings.UTF Encoding;
3
   use Ada.Strings.UTF Encoding;
4
5
   with Ada.Strings.UTF Encoding.Wide Wide Strings;
6
   use Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
7
8
   procedure Show_UTF_8 is
9
10
```

```
Symbol_UTF_8 : constant UTF_8_String := "x";
11
      Symbol_WWS : constant Wide_Wide_String :=
12
                         Decode (Symbol_UTF_8);
13
14
   begin
15
      Put_Line ("Wide_Wide_String Length: "
16
                 & Symbol_WWS'Length'Image);
17
      Put_Line ("Wide_Wide_String Size:
18
                 & Symbol_WWS'Size'Image);
19
      Put_Line ("UTF-8 String Length:
20
                 & Symbol_UTF_8'Length'Image);
21
      Put_Line ("UTF-8 String Size:
22
                 & Symbol_UTF_8'Size'Image);
23
      New_Line;
24
      Put_Line ("UTF-8 String:
                                              n.
25
                 & Symbol_UTF_8);
26
   end Show_UTF_8;
27
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8
MD5: 67653dfd377f04b32421cf09b25939fe

Runtime output

Wide_Wide_String Length: 1 Wide_Wide_String Size: 32 UTF-8 String Length: 4 UTF-8 String Size: 32 UTF-8 String: x

In this case, both strings — using the Wide_Wide_String type or the UTF_8_String type — have the same size: 32 bits. (Here, we're using the x symbol from the Mathematical Alphanumeric Symbols $block^{143}$, not the standard "x" from the Basic Latin $block^{144}$.)

7.3.3 UTF-16 encoding and decoding

So far, we've discussed the UTF-8 encoding scheme. However, other encoding schemes exist and are supported as well. In fact, the Ada.Strings.UTF_Encoding package defines three encoding schemes:

```
type Encoding_Scheme is (UTF_8,
UTF_16BE,
UTF_16LE);
```

For example, instead of using UTF-8 encoding, we can use UTF-16 encoding — either in the big-endian or in the little-endian version. To convert between UTF-8 and UTF-16 encoding schemes, we can make use of the conversion functions from the Ada.Strings. UTF_Encoding.Conversions package.

To declare a UTF-16 encoded string, we can use one of the following data types:

- the 8-bit-character based UTF_String type, or
- the 16-bit-character based UTF_16_Wide_String type.

When using the 8-bit version, though, we have to specify the input and output schemes when converting between UTF-8 and UTF-16 encoding schemes.

¹⁴³ https://en.wikipedia.org/wiki/Mathematical_Alphanumeric_Symbols

¹⁴⁴ https://en.wikipedia.org/wiki/Basic_Latin_(Unicode_block)

Let's see a code example that makes use of both UTF_String and UTF_16_Wide_String types:

```
Listing 11: show_utf16_types.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Ada.Strings.UTF Encoding;
3
   use Ada.Strings.UTF Encoding;
4
5
   with Ada.Strings.UTF Encoding.Conversions;
6
   use Ada.Strings.UTF_Encoding.Conversions;
7
8
   procedure Show_UTF16_Types is
9
       Symbols_UTF_8 : constant
10
         UTF_8_String := "♥♪";
11
12
       Symbols UTF 16 : constant
13
         UTF 16 Wide String :=
14
           Convert (Symbols_UTF_8);
15
           ^ Calling Convert for UTF_8_String
16
             to UTF_16_Wide_String conversion.
       - -
17
18
       Symbols UTF 16BE : constant
19
         UTF_String :=
20
           Convert (Item
                                    => Symbols_UTF_8,
21
                     Input Scheme => UTF 8,
22
                     Output Scheme => UTF 16BE);
23
          ^ Calling Convert for UTF 8 String
24
             to UTF String conversion in UTF-16BE
       - -
25
       - -
             encoding.
26
   begin
27
       Put_Line ("UTF_8_String:
                                            ш
28
                 & Symbols_UTF_8);
29
30
      Put_Line ("UTF_16_Wide_String:
                                            н
31
                 & Convert (Symbols_UTF_16));
32
                     Calling Convert for
33
                      the UTF_16_Wide_String to
34
                      UTF 8 String conversion.
35
36
      Put Line
37
         ("UTF_String / UTF_16BE: "
38
          & Convert
39
                              => Symbols_UTF_16BE,
              (Item
40
               Input Scheme => UTF 16BE,
41
               Output_Scheme => UTF_8));
42
   end Show_UTF16_Types;
43
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_16_Types
MD5: 905e20e83a6199fdc91a6b15bb71bb01

Runtime output

UTF_8_String:	♥Ĵ
UTF_16_Wide_String:	♥♪]
<pre>UTF_String / UTF_16BE:</pre>	۷Л

In this example, we're declaring a UTF-8 encoded string and storing it in the Symbols_UTF_8 constant. Then, we're calling the Convert functions to convert between UTF-8 and UTF-16 encoding schemes. We're using two versions of this function:

- the Convert function that returns an object of UTF_16_Wide_String type for an input of UTF_8_String type, and
- the Convert function that returns an object of UTF_String type for an input of UTF_8_String type.
 - In this case, we need to specify the input and output schemes (see Input_Scheme and Output_Scheme parameters in the code example).

Previously, we've seen that the Ada.Strings.UTF_Encoding.Wide_Wide_Strings package offers functions to convert between UTF-8 and the Wide_Wide_String type. The same kind of conversion functions exist for UTF-16 strings as well. Let's look at this code example:

Listing 12: show_ww_utf16_string.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Strings.UTF Encoding;
3
   use Ada.Strings.UTF_Encoding;
4
5
   with Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
6
   use Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
7
8
   with Ada.Strings.UTF_Encoding.Conversions;
9
   use Ada.Strings.UTF_Encoding.Conversions;
10
11
   procedure Show_WW_UTF16_String is
12
       Symbols_UTF_16 : constant
13
        UTF 16_Wide_String :=
14
           Wide_Character'Val (16#2665#) &
15
           Wide_Character'Val (16#266B#);
16
          ^ Calling Wide_Character'Val
17
             to specify the UTF-16 BE code
18
       - -
       - -
             19
20
      Symbols_WWS : constant
21
        Wide_Wide_String :=
22
           Decode (Symbols_UTF_16);
23
           ^ Calling Decode for UTF_16_Wide String
24
       - -
             to Wide_Wide_String conversion.
25
   begin
26
      Put_Line ("UTF_16_Wide_String: "
27
                 & Convert (Symbols UTF 16));
28
                   ^ Calling Convert for the
29
       - -
                     UTF_16_Wide_String to
       - -
30
                     UTF_8_String conversion.
       - -
31
32
      Put_Line ("Wide_Wide_String:
33
                 & Encode (Symbols_WWS));
34
                    Calling Encode for the
35
                     Wide Wide String to
36
                     UTF_8_String conversion.
37
   end Show_WW_UTF16_String;
38
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.WW_UTF_16_String MD5: 900af8f5c6aad7303c3e49c1c4a68d73

Runtime output

UTF_16_Wide_String: ♥♪ Wide_Wide_String: ♥♪ In this example, we're calling the Wide_Character'Val function to specify the UTF-16 BE code of the "♥" and "♪" symbols. We're then using the Decode function to convert between the UTF_16_Wide_String and the Wide_Wide_String types.

7.4 UTF-8 applications

In this section, we take a further look into UTF-8 encoding and some real-world applications. First, we discuss the use of UTF-8 encoding in source-code files. Then, we talk about parsing UTF-8 files using *wide-wide* strings.

7.4.1 UTF-8 encoding in source-code files

In the past, it was common to use different character sets in text files when writing in different (human) languages. By default, Ada source-code files are expected to use the Latin-1 coding, which is a 8-bit character set.

Nowadays, however, using UTF-8 coding for text files — including source-code files — is very common. If your Ada code only uses standard ASCII characters, but you're saving it in a UTF-8 coded file, there's no need to worry about character sets, as UTF-8 is backwards compatible with ASCII.

However, you might want to use Unicode symbols in your Ada source code to declare constants — as we did in the previous sections — and store the source code in a UTF-8 coded file. In this case, you need be careful about how this file is parsed by the compiler.

Let's look at this source-code example:

```
Listing 13: show_utf_8_strings.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Strings.UTF_Encoding;
3
   use Ada.Strings.UTF_Encoding;
4
5
   procedure Show_UTF_8_Strings is
6
7
        Symbols UTF 8 : constant
8
          UTF_8_String := "♥♪";
9
10
   begin
11
        Put_Line ("UTF_8_String: "
12
                  & Symbols_UTF_8);
13
14
        Put_Line ("Length:
                                   н
15
                  & Symbols_UTF_8'Length'Image);
16
17
   end Show_UTF_8_Strings;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8_Strings
MD5: fd1aaff161a33365d15adca5bea7b277
```

Runtime output

UTF_8_String: ♥♪ Length: 6

Here, we're using Unicode symbols to initialize the Symbols_UTF_8 constant of UTF_8_String type.

Now, let's assume this source-code example is stored in a UTF-8 coded file. Because the "•]" string makes use of non-ASCII Unicode symbols, representing this string in UTF-8 format will require more than 2 bytes. In fact, each one of those Unicode symbols requires 2 bytes to be encoded in UTF-8. (Keep in mind that Unicode symbols may require between 1 to 4 bytes¹⁴⁵ to be encoded in UTF-8 format.) Also, in this case, the UTF-8 encoding process is using two additional bytes. Therefore, the total length of the string is six, which matches what we see when running the Show UTF 8 Strings procedure. In other words, the length of the Symbols UTF 8 string doesn't refer to those two characters (" \P]") that we were using in the constant declaration, but the length of the encoded bytes in its UTF-8 representation.

The UTF-8 format is very useful for storing and transmitting texts. However, if we want to process Unicode symbols, it's probably better to use string types with 32-bit characters such as Wide Wide String. For example, let's say we want to use the "*]" string again to initialize a constant of Wide Wide String type:

Listing 14: show wws strings.adb

```
with Ada.Text_I0;
1
   with Ada.Wide_Wide_Text_I0;
2
З
   procedure Show WWS Strings is
4
5
      package TIO
                     renames Ada.Text I0;
6
7
      package WWTIO renames Ada.Wide Wide Text IO;
8
      Symbols WWS : constant
9
        Wide Wide String := "♥♪";
10
11
   beain
12
      WWTIO.Put_Line ("Wide_Wide_String: "
13
                        & Symbols WWS);
14
15
      TIO.Put Line ("Length:
16
                      & Symbols WWS'Length'Image);
17
18
   end Show_WWS_Strings;
```

Code block metadata

Project: Courses.Advanced Ada.Data Types.Strings.String Encoding.WWS Strings W8 MD5: 1e5e38e62b412de48d3fa4271bb48bf1

Runtime output

19

Wide Wide String: ♥♪ Length:

In this case, as mentioned above, if we store this source code in a text file using UTF-8 format, we need to ensure that the UTF-8 coded symbols are correctly interpreted by the compiler when it parses the text file. Otherwise, we might get unexpected behavior. (Interpreting the characters in UTF-8 format as Latin-1 format is certainly an example of what we want to avoid here.)

In the GNAT toolchain

You can use UTF-8 coding in your source-code file and initialize strings of 32-bit characters. However, as we just mentioned, you need to make sure that the UTF-8 coded symbols are correctly interpreted by the compiler when dealing with types such as

¹⁴⁵ https://en.wikipedia.org/wiki/UTF-8

Wide Wide String. For this case, GNAT offers the -gnatW8 switch. Let's run the previous example using this switch: Listing 15: show wws strings.adb with Ada.Text IO; with Ada.Wide_Wide_Text_I0; procedure Show_WWS_Strings is package TIO renames Ada.Text I0; package WWTI0 renames Ada.Wide_Wide_Text_I0; Symbols WWS : constant Wide_Wide_String := "♥♪"; begin WWTIO.Put_Line ("Wide_Wide_String: " & Symbols_WWS); TIO.Put_Line ("Length: & Symbols WWS'Length'Image); end Show WWS Strings; Code block metadata Project: Courses.Advanced Ada.Data Types.Strings.String Encoding.WWS Strings W8 MD5: 1e5e38e62b412de48d3fa4271bb48bf1 Runtime output Wide Wide String: ♥♪ Length: 2 Because the Wide Wide String type has 32-bit characters. we expect the length of the string to match the number of symbols that we're using. Indeed, when running the Show WWS Strings procedure, we see that the Symbols WWS string has a length of two characters, which matches the number of characters of the " When we use the -gnatW8 switch, GNAT converts the UTF-8-coded string ("*]) to UTF-32 format, so we get two 32-bit characters. It then uses the UTF-32-coded string to initialize the Symbols WWS string. If we don't use the -gnatW8 switch, however, we get wrong results. Let's look at the same example again without the switch: Listing 16: show wws strings.adb

```
with Ada.Text IO;
1
   with Ada.Wide_Wide_Text_I0;
2
3
   procedure Show_WWS_Strings is
4
5
       package TIO
                      renames Ada.Text I0;
6
       package WWTIO renames Ada.Wide_Wide_Text_I0;
7
8
9
       Symbols WWS : constant
10
         Wide_Wide_String := "♥♪";
11
   begin
12
       WWTIO.Put_Line ("Wide_Wide_String: "
13
                        & Symbols WWS);
14
15
       TIO.Put Line ("Length:
16
```

1 2

3

4 5

6

7 8

g

10 11 12

13

14 15

16

17 18

19

17 18 & Symbols_WWS'Length'Image);

19 end Show_WWS_Strings;

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.WWS_Strings_No_ ⇔W8

MD5: 1e5e38e62b412de48d3fa4271bb48bf1

Runtime output

Wide_Wide_String: ♥♪ Length: 6

Now, the "•]" string is being interpreted as a string of six 8-bit characters. (In other words, the UTF-8-coded string isn't converted to the UTF-32 format.) Each of those 8-bit characters is then stored in a 32-bit character of the Wide_Wide_String type. This explains why the Show_WWS_Strings procedure reports a length of 6 components for the Symbols_WWS string.

Portability of UTF-8 in source-code files

In a previous code example, we were assuming that the format that we use for the sourcecode file is UTF-8. This allows us to simply use Unicode symbols directly in strings:

Symbol_UTF_8 : constant UTF_8_String := "*";

This approach, however, might not be portable. For example, if the compiler uses a different string encoding for source-code files, it might interpret that Unicode character as something else — or just throw a compilation error.

If you're afraid that format mismatches might happen in your compilation environment, you may want to write strings in your code in a completely portable fashion, which consists in entering the exact sequence of codes in bytes — using the **Character'Val** function — for the symbols you want to use.

We can reuse parts of the previous example and replace the UTF-8 character with the corresponding UTF-8 code:

Listing	17:	show	utf	8.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Ada.Strings.UTF_Encoding;
3
   use Ada.Strings.UTF_Encoding;
4
5
   procedure Show_UTF_8 is
6
7
      Symbol_UTF_8 : constant
8
        UTF_8_String :=
9
           Character'Val (16#e2#)
10
           & Character'Val (16#98#)
11
           & Character'Val (16#85#);
12
13
   begin
14
      Put_Line ("UTF-8 String: "
15
                 & Symbol_UTF_8);
16
   end Show_UTF_8;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8
MD5: 8ff02bc1793c0c5ac1ff24f62941af73
```

Runtime output

```
UTF-8 String: *
```

Here, we use a sequence of three calls to the Character'Val(code) function for the UTF-8 code that corresponds to the " \star " symbol.

7.4.2 Parsing UTF-8 files for Wide-Wide-String processing

A typical use-case is to parse a text file in UTF-8 format and use *wide-wide* strings to process the lines of that file. Before we look at the implementation that does that, let's first write a procedure that generate a text file in UTF-8 format:

```
Listing 18: generate_utf_8_file.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Strings.UTF_Encoding;
3
   use Ada.Strings.UTF_Encoding;
4
5
   procedure Generate UTF 8 File
6
     (Output File Name : String)
7
8
   is
      F : File_Type;
9
10
   begin
      Create (F, Out File, Output File Name);
11
      Put_Line (F, UTF_8_String'("♥♪"));
12
      Put Line
13
        (F,
14
         UTF_8_String'((عالم" يا مرحبا");
15
      Close (F);
16
   end Generate UTF 8 File;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8_File_

←Processing

MD5: 58c7591796bc1348796afa6db6f64d22
```

Procedure Generate_UTF_8_File writes two strings with non-Latin characters into the UTF-8 file indicated by the Output_File_Name parameter.

In addition, let's implement an auxiliary procedure to display the individual characters of a *wide-wide* string:

Listing 19: put_line_utf_8_characters.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Strings.UTF Encoding;
3
   use Ada.Strings.UTF Encoding;
4
   with Ada.Strings.UTF Encoding.Wide Wide Strings;
6
   use Ada.Strings.UTF Encoding.Wide Wide Strings;
7
8
   procedure Put_Line_UTF_8_Characters
9
     (WSS : Wide_Wide_String)
10
   is
11
```

```
procedure Put_Complete_UTF_8_String
12
         (WSS : Wide_Wide_String)
13
       is
14
          S_UTF_8 : constant UTF_8_String :=
15
                       Encode (WSS);
16
      begin
17
          Put_Line ("STRING: " & S_UTF_8);
18
          Put_Line ("Length: "
19
                    & WSS'Length'Image
20
                     & " characters");
21
          New Line;
22
      end Put_Complete_UTF_8_String;
23
24
       -- This is a wrapper function of the
25
       -- Encode function for the
26
       -- Wide_Wide_Character type:
27
      function Encode (Item : Wide_Wide_Character)
28
                         return UTF_8_String
29
        is
30
           SC : constant Wide_Wide_String (1 .. 1)
31
                  := (1 => Item);
32
           -- We need a 1-character string
33
           -- for the call to Encode.
34
       begin
35
           return Encode (SC);
36
37
      end Encode;
38
      procedure Put_UTF_8_Characters
39
         (WSS : Wide_Wide_String) is
40
      begin
41
          for I in WSS'Range loop
42
             Put (I'Image & ": ");
43
             Put (Encode (WSS (I)));
44
             New_Line;
45
          end loop;
46
      end Put_UTF_8_Characters;
47
48
   begin
49
        Put_Complete_UTF_8_String (WSS);
50
        Put_UTF_8_Characters (WSS);
51
        Put Line ("--
                                       ---");
52
   end Put_Line_UTF_8_Characters;
53
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8_File_ ⊶Processing MD5: 14fae1f2b1d3795f3cef244f60082fcc

Finally, let's look at a code example that parses an UTF-8 file:

Listing 20: show utf 8.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Strings.UTF_Encoding;
duse Ada.Strings.UTF_Encoding;
with Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
duse Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
```

```
with Generate_UTF_8_File;
9
   with Put_Line_UTF_8_Characters;
10
11
   procedure Show_UTF_8 is
12
13
       File_Name : constant String :=
14
                      "utf-8_test.txt";
15
16
       procedure Read_UTF_8_File
17
         (Input_File_Name : String)
18
       is
19
          F : File_Type;
20
21
       begin
          Open (F, In_File, Input_File_Name);
22
23
          while not End_Of_File (F) loop
24
             declare
25
                S_UTF8 : constant UTF_8_String
26
                             := Get_Line (F);
27
                 S
                        : constant Wide_Wide_String
28
                             := Decode (S_UTF8);
29
             begin
30
                Put_Line_UTF_8_Characters (S);
31
             end;
32
33
          end loop;
34
          Close (F);
       end Read_UTF_8_File;
35
36
   begin
37
       Generate_UTF_8_File (File_Name);
38
       Read_UTF_8_File (File_Name);
39
   end Show_UTF_8;
40
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8_File_

⊶Processing

MD5: 512ad5ac7c6d5936735f017bfe629aa3
```

Runtime output

```
STRING: ♥♪
Length: 2 characters
1: ♥
2: រា
عالم يا مرحبا :STRING
Length: 13 characters
م :1
 ر :2
 3: כ
 ب :4
 5: I
 6:
 ي :7
ا :8
 9:
 ع :10
 11: |
```

```
ل :12
م :13
```

The Show_UTF_8 procedure first calls the Generate_UTF_8_File procedure to generate a text file in UTF-8 format, and then calls the nested Read_UTF_8_File procedure to read from that file — this is done by reading the 8-bit UTF-8 encoded string and decoding it into a string of Wide_Wide_String type.

(Note that we call the auxiliary Put_Line_UTF_8_Characters procedure to display the characters of each line we read from the UTF-8 file.)

For completeness, we include the nested Read_Write_UTF_8_File procedure, which not only reads each line from a UTF-8 file, but also writes it into another UTF-8 file:

Listing 21: show_utf_8.adb

```
1
   with Ada.Text_IO; use Ada.Text_IO;
2
   with Ada.Strings.UTF_Encoding;
3
   use Ada.Strings.UTF_Encoding;
4
5
   with Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
6
   use Ada.Strings.UTF_Encoding.Wide_Wide_Strings;
7
8
   with Generate_UTF_8_File;
9
   with Put_Line_UTF_8_Characters;
10
11
   procedure Show_UTF_8 is
12
13
       File_Name_In : constant String :=
14
                          "utf-8 test.txt";
15
       File Name Out : constant String :=
16
                          "utf-8_copy.txt";
17
18
       procedure Read_Write_UTF_8_File
19
20
         (Input_File_Name,
         Output_File_Name : String)
21
      is
22
          F_In, F_Out : File_Type;
23
      begin
24
          Open (F_In, In_File, Input_File_Name);
25
          Create (F_Out, Out_File, Output_File_Name);
26
27
28
          while not End_Of_File (F_In) loop
             declare
29
                S : constant Wide_Wide_String :=
30
                       Decode (Get_Line (F_In));
31
             begin
32
                Put_Line_UTF_8_Characters (S);
33
                Put_Line (F_Out, Encode (S));
34
             end;
35
          end loop;
36
37
          Close (F_In);
38
39
          Close (F_Out);
       end Read_Write_UTF_8_File;
40
41
   begin
42
       Generate_UTF_8_File (File_Name_In);
43
44
```

```
45 Read_Write_UTF_8_File
46 (Input_File_Name => File_Name_In,
47 Output_File_Name => File_Name_Out);
48 end Show_UTF_8;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8_File_

⊶Processing

MD5: 8cd13e8a565266fa5dd854ff6a34524c
```

Runtime output

```
STRING: ♥♪
Length: 2 characters
1: ♥
2: Л
عالم يا مرحبا :STRING
Length: 13 characters
م :1
ر :2
3: כ
ب :4
5: I
6:
ي :7
8: Ï
9:
ع :10
11: |
12: J
م :13
```

In the nested Read_Write_UTF_8_File procedure, we see both Decode and Encode functions being called to convert from and to the UTF_8_String type, respectively.

In the GNAT toolchain

If we use the -gnatW8 switch, which we mentioned *in a previous section* (page 328), the implementation of Generate_UTF_8_File and Put_Line_UTF_8_Characters must be adapted. In addition, we can simplify the implementation of the Show_UTF_8 procedure, too. (Note, however, that the previous implementation, which makes use of the Decode and Encode functions, would work fine as well.)

Listing 22: put_line_utf_8_characters.adb

```
with Ada.Wide Wide Text I0;
1
   use Ada.Wide_Wide_Text_I0;
2
3
   procedure Put Line UTF 8 Characters
4
5
     (WSS : Wide_Wide_String)
6
   is
      procedure Put_Complete_UTF_8_String
7
        (WSS : Wide_Wide_String)
8
      is
9
      begin
10
          Put_Line ("STRING: " & WSS);
11
```

```
Put Line ("Length: "
12
                     & WSS'Length'Wide_Wide_Image
13
                     & " characters");
14
          New Line;
15
       end Put_Complete_UTF_8_String;
16
17
       procedure Put_UTF_8_Characters
18
         (WSS : Wide Wide String)
19
       is
20
       begin
21
          for I in WSS'Range loop
22
             Put (I'Wide_Wide_Image & ": ");
23
             Put (WSS (I));
24
             New Line;
25
          end loop;
26
       end Put_UTF_8_Characters;
27
28
   begin
29
        Put Complete UTF 8 String (WSS);
30
        Put_UTF_8_Characters (WSS);
31
        Put Line ("--
                                       ---");
32
   end Put Line UTF 8 Characters;
33
                              Listing 23: generate_utf_8_file.adb
   with Ada.Wide_Wide_Text_I0;
1
2
   use Ada.Wide_Wide_Text_I0;
3
   procedure Generate_UTF_8_File
4
      (Output_File_Name : String)
5
   is
6
7
       F : File_Type;
   begin
8
       Create (F, Out_File, Output_File_Name);
9
       Put_Line (F, "♥♪");
10
       Put_Line (F, "عالم" يا مرحبا);
11
12
       Close (F);
   end Generate_UTF_8_File;
13
                                  Listing 24: show utf 8.adb
   with Ada.Wide_Wide_Text_I0;
1
   use Ada.Wide_Wide_Text_I0;
2
3
   with Generate_UTF_8_File;
4
   with Put_Line_UTF_8_Characters;
5
6
   procedure Show_UTF_8 is
7
8
       File_Name_In : constant String :=
9
                          "utf-8_test.txt";
10
       File_Name_Out : constant String :=
11
12
                          "utf-8_copy.txt";
13
       procedure Read_Write_UTF_8_File
14
         (Input_File_Name,
15
          Output_File_Name : String)
16
17
       is
          F_In, F_Out : File_Type;
18
       begin
19
          Open (F_In, In_File, Input_File_Name);
20
          Create (F_Out, Out_File, Output_File_Name);
21
22
```

```
while not End_Of_File (F_In) loop
23
             declare
24
                S : constant Wide_Wide_String :=
25
                       Get_Line (F_In);
26
             begin
27
                Put_Line_UTF_8_Characters (S);
28
                Put_Line (F_Out, S);
29
             end:
30
          end loop;
31
32
          Close (F_In);
Close (F_Out);
33
34
       end Read_Write_UTF_8_File;
35
36
   begin
37
      Generate_UTF_8_File (File_Name_In);
38
39
       Read_Write_UTF_8_File
40
         (Input File Name => File Name In,
41
          Output File Name => File Name Out);
42
   end Show_UTF_8;
43
   Code block metadata
   Project: Courses.Advanced_Ada.Data_Types.Strings.String_Encoding.UTF_8_File_
    →Processing
   MD5: 8eeed924f6d661a0a62ecb4d94be7027
   Runtime output
   STRING: ♥♪
   Length: 2 characters
    1: ♥
    2: ♪
   عالم يا مرحبا :STRING
   Length: 13 characters
    م :1
    ر :2
    3: ෭
    ب :4
    5: I
    6:
    ي :7
    8: I
    9:
    ع :10
    11: |
    12: J
    م :13
            - - - - - - - - - - - - -
   In this version of the code, we've removed all references to the UTF 8 String type
   - as well as the Decode and Encode functions that we were using to convert from
   and to this type. In this case, all UTF-8 processing happens directly using strings of
   Wide Wide Strings type.
```

7.5 Image attribute

7.5.1 Overview

In the Introduction to Ada^{146} course, we've seen that the Image attribute returns a string that contains a textual representation of an object. For example, we write **Integer**'Image (V) to get a string for the integer variable V:

Listing 25: show simple image.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Show_Simple_Image is
V : Integer;
begin
V := 10;
Put_Line ("V: " & Integer'Image (V));
end Show_Simple_Image;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Simple_Image
MD5: e38f6f1a0808f12bd53c1f3cf4983353
```

Runtime output

V: 10

Naturally, we can use the Image attribute with other scalar types. For example:

Listing 26: show_simple_image.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
    procedure Show_Simple_Image is
3
       type Status is (Unknown, Off, On);
4
5
       V : Float;
6
       S : Status;
7
   begin
8
       V := 10.0;
9
       S := Unknown;
10
11
       Put_Line ("V: " & Float'Image (V));
Put_Line ("S: " & Status'Image (S));
12
13
   end Show_Simple_Image;
14
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Simple_Image
MD5: d3369518b610b7bf6c8dcefdecdb0c44

Runtime output

```
V: 1.00000E+01
S: UNKNOWN
```

In this example, we retrieve a string representing the floating-point variable V. Also, we use Status'Image (V) to retrieve a string representing the textual version of the Status.

 $^{146}\ https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html# intro-ada-image-attribute$

In the Ada Reference Manual

```
    Image Attributes<sup>147</sup>
```

7.5.2 Type'Image and Obj'Image

We can also apply the Image attribute to an object directly:

Listing 27: show_simple_image.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show Simple Image is
3
      V : Integer;
4
5
   begin
      V := 10;
6
      Put_Line ("V: " & V'Image);
7
8
      -- Equivalent to:
9
          Put_Line ("V: " & Integer'Image (V));
      - -
10
   end Show Simple Image;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Simple_Image
MD5: c8b2e458de47b403568dd795b3d3fc24

Runtime output

V: 10

In this example, the **Integer** 'Image (V) and V'Image forms are equivalent.

7.5.3 Wider versions of Image

Although we've been talking only about the Image attribute, it's important to mention that each of the wider versions of the string types also has a corresponding Image attribute. In fact, this is the attribute for each string type:

Attribute	Type of Returned String
Image	String Wide String
Wide_Image	Wide_String
Wide_Wide_Image	Wide_Wide_String

Let's see a simple example:

Listing 28: show wide wide image.adb

```
1 with Ada.Wide_Wide_Text_I0;
2 use Ada.Wide_Wide_Text_I0;
3 
4 procedure Show_Wide_Wide_Image is
5 F : Float;
6 begin
7 F := 100.0;
```

(continues on next page)

¹⁴⁷ http://www.ada-auth.org/standards/22rm/html/RM-4-10.html

```
8 Put_Line ("F = "
9 & F'Wide_Wide_Image);
10 end Show_Wide_Wide_Image;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Wide_Wide_Image
MD5: ff542ef93286529343466c27935d5c21
```

Runtime output

F = 1.00000E+02

In this example, we use the Wide_Wide_Image attribute to retrieve a string of Wide_Wide_String type for the floating-point variable F.

7.5.4 Image attribute for non-scalar types

Note

This feature was introduced in Ada 2022.

In the previous code examples, we were using the Image attribute with scalar types, but it isn't restricted to those types. In fact, we can also use this attribute when dealing with non-scalar types. For example:

	Listing	29:	simple	records.ads
--	---------	-----	--------	-------------

```
package Simple_Records is
1
2
       type Rec is limited private;
3
4
       type Rec_Access is access Rec;
5
6
       function Init return Rec;
7
8
       type Null_Rec is null record;
9
10
   private
11
12
       type Rec is limited record
13
          F : Float;
14
          I : Integer;
15
       end record;
16
17
       function Init return Rec is
18
          ((F \implies 10.0, I \implies 4));
19
20
   end Simple_Records;
21
```

Listing 30: show_non_scalar_image.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Unchecked_Deallocation;
with Simple_Records;
use Simple_Records;
```

```
procedure Show_Non_Scalar_Image is
7
8
      procedure Free is
9
         new Ada.Unchecked_Deallocation
10
           (Object => Rec,
11
            Name => Rec_Access);
12
13
      R_A : Rec_Access :=
14
        new Rec'(Init);
15
16
      N R : Null Rec :=
17
         (null record);
18
19
   begin
      R_A := new Rec'(Init);
20
      N_R := (null record);
21
22
                            " & R A'Image);
      Put Line ("R A:
23
      Put_Line ("R_A.all: " & R_A.all'Image);
24
      Put_Line ("N_R:
                            " & N_R'Image);
25
26
      Free (R A);
27
      Put_Line ("R_A:
                            " & R_A'Image);
28
   end Show_Non_Scalar_Image;
29
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Non_Scalar_Image
MD5: eb48f3fbe69b70258bc26f467918717c

Runtime output

R_A: (access 5f2a5ad892c0)
R_A.all:
(F => 1.00000E+01,
I => 4)
N_R: (NULL RECORD)
R_A: null

In the Show_Non_Scalar_Image procedure from this example, we display the access value of R_A and the contents of the dereferenced access object (R_A.all). Also, we see the indication that N_R is a null record and R_A is null after the call to Free.

Historically

Since Ada 2022, the Image attribute is available for all types. Prior to this version of the language, it was only available for scalar types. (For other kind of types, programmers had to use the Image attribute for each component of a record, for example.)

In fact, prior to Ada 2022, the Image attribute was described in the 3.5 Scalar Types¹⁴⁸ section of the Ada Reference Manual, as it was only applied to those types. Now, it is part of the new Image Attributes¹⁴⁹ section.

Let's see another example, this time with arrays:

¹⁴⁸ http://www.ada-auth.org/standards/22rm/html/RM-3-5.html

¹⁴⁹ http://www.ada-auth.org/standards/22rm/html/RM-4-10.html

Listing 31: show_array_image.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show_Array_Image is
3
4
      type Float_Array is
5
        array (Positive range <>) of Float;
6
7
      FA_3C : Float_Array (1 .. 3);
8
      FA Null : Float Array (1 .. 0);
9
10
   begin
11
      FA_3C := [1.0, 3.0, 2.0];
12
      FA_Null := [];
13
14
      Put Line ("FA 3C: " & FA 3C'Image);
15
      Put_Line ("FA_Null: " & FA_Null'Image);
16
   end Show Array Image;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Array_Image
MD5: a24daba1d92a139ae8995bba5a81e0d6

Runtime output

```
FA_3C:
[ 1.00000E+00, 3.00000E+00, 2.00000E+00]
FA_Null:
[]
```

In this example, we display the values of the three components of the FA_3C array. Also, we display the null array FA_Null.

7.5.5 Image attribute for tagged types

In addition to untagged types, we can also use the Image attribute with tagged types. For example:

Listing 32: simple_records.ads

```
package Simple Records is
1
2
      type Rec is tagged limited private;
3
4
      function Init return Rec;
5
6
      type Rec_Child is new Rec with private;
7
8
      overriding function Init return Rec_Child;
9
10
   private
11
12
      type Status is (Unknown, Off, On);
13
14
       type Rec is tagged limited record
15
          F : Float;
16
          I : Integer;
17
      end record;
18
19
```

```
function Init return Rec is
20
          ((F \implies 10.0, I \implies 4));
21
22
       type Rec_Child is new Rec with record
23
          Z : Status;
24
       end record;
25
26
       function Init return Rec Child is
27
           (Rec'(Init) with Z => Off);
28
29
   end Simple_Records;
```

Listing 33: show_tagged_image.adb

```
with Ada.Text IO;
                        use Ada.Text I0;
1
2
   with Simple Records; use Simple Records;
3
4
   procedure Show_Tagged_Image is
5
      R : constant Rec
                                    := Init;
6
      R_Class : constant Rec'Class := Rec'(Init);
7
      R_C : constant Rec_Child := Init;
8
   begin
9
                          " & R'Image);
      Put_Line ("R:
10
      Put_Line ("R_Class: " & R_Class'Image);
11
                          " & R_C'Image);
      Put Line ("R A:
12
  end Show_Tagged_Image;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Tagged_Image MD5: 496827d5f81f8b7bec3b1d4a104f550e

Runtime output

30

```
R:
(F => 1.00000E+01,
I => 4)
R Class: SIMPLE RECORDS.REC'
(F => 1.00000E+01,
I => 4)
R_A:
(F => 1.00000E+01,
I => 4,
 Z \implies OFF)
```

In the Show Tagged Image procedure from this example, we display the contents of the R object of Rec type and the R Class object of Rec 'Class type. Also, we display the contents of the R C object of the Rec Child type, which is derived from the Rec type.

7.5.6 Image attribute for task and protected types

We can also apply the Image attribute to protected objects and tasks:

Listing 34: simple_tasking.ads

```
package Simple Tasking is
1
2
      protected type Protected_Float (I : Integer) is
3
Δ
```

```
private
5
         V : Float := Float (I);
6
      end Protected_Float;
7
8
      protected type Protected_Null is
9
10
      private
      end Protected_Null;
11
12
      task type T is
13
         entry Start;
14
      end T;
15
16
   end Simple_Tasking;
17
```

Listing 35: simple_tasking.adb

```
package body Simple Tasking is
1
2
      protected body Protected_Float is
3
4
      end Protected_Float;
5
6
      protected body Protected_Null is
7
8
      end Protected Null;
9
10
      task body T is
11
      begin
12
          accept Start;
13
      end T;
14
15
   end Simple_Tasking;
16
```

Listing 36: show_protected_task_image.adb

```
use Ada.Text_I0;
   with Ada.Text_I0;
1
2
   with Simple_Tasking; use Simple_Tasking;
3
4
   procedure Show_Protected_Task_Image is
5
6
      PF : Protected Float (0);
7
8
      PN : Protected Null;
9
      T1 : T;
10
11
   begin
      Put_Line ("PF: " & PF'Image);
12
      Put_Line ("PN: " & PN'Image);
13
      Put_Line ("T1: " & T1'Image);
14
15
      T1.Start;
16
   end Show_Protected_Task_Image;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Image_Attribute.Protected_Task_

GImage

MD5: feb14f17ba1cca0311420272ef91ab38
```

Runtime output

```
PF: (protected object)
PN: (protected object)
T1: (task t1_00005CCB29E40090)
```

In this example, we display information about the protected object PF, the componentless protected object PN and the task T1.

7.6 Put_Image aspect

\rm Note

This feature was introduced in Ada 2022.

7.6.1 Overview

In the previous section, we discussed many details about the Image attribute. In the code examples from that section, we've seen the default behavior of this attribute: the string returned by the calls to Image was always in the format defined by the Ada standard.

In some situations, however, we might want to customize the string that is returned by the Image attribute of a type T. Ada allows us to do that via the Put_Image aspect. This is what we have to do:

- 1. Specify the Put_Image aspect for the type T and indicate a procedure with a specific parameter profile let's say, for example, a procedure named P.
- 2. Implement the procedure P and write the information we want to use into a buffer (by calling the routines defined for Root_Buffer_Type, such as the Put procedure).

We can see these steps performed in the code example below:

```
Listing 37: show put image.ads
```

```
with Ada.Strings.Text_Buffers;
1
2
   package Show Put Image is
3
4
      type T is null record
5
        with Put_Image => Put_Image_T;
6
              ^ Custom version of Put Image
       - -
7
8
      use Ada.Strings.Text_Buffers;
9
10
      procedure Put_Image_T
11
         (Buffer : in out Root_Buffer_Type'Class;
12
                           T);
         Arg
                 1
13
14
   end Show_Put_Image;
15
```

Listing 38: show_put_image.adb

```
package body Show Put Image is
1
2
      procedure Put Image T
3
        (Buffer : in out Root_Buffer_Type'Class;
4
                         T) is
         Ara
5
               .
         pragma Unreferenced (Arg);
6
      begin
7
             Call Put with customized
8
```

```
9 -- information
10 Buffer.Put ("<custom info>");
11 end Put_Image_T;
12
13 end Show_Put_Image;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Put_Image.Simple_Put_Image
MD5: 45c55444f0e1825312b5eafe307ca58d

In the Show_Put_Image package, we use the Put_Image aspect in the declaration of the T type. There, we indicate that the Image attribute shall use the Put_Image_T procedure instead of the default version.

In the body of the Put_Image_T procedure, we implement our custom version of the Image attribute. We do that by calling the Put procedure with the information we want to provide in the Image attribute. Here, we access a buffer of Root_Buffer_Type type, which is defined in the Ada.Strings.Text_Buffers package. (We discuss more about this package *later on* (page 352).)

1 In the Ada Reference Manual

```
• Image Attributes<sup>150</sup>
```

7.6.2 Complete Example of Put_Image

Let's see a complete example in which we use the Put_Image aspect and write useful information to the buffer:

```
Listing 39: custom_numerics.ads
```

```
with Ada.Strings.Text_Buffers;
1
2
   package Custom Numerics is
3
4
      type Float_Integer is record
5
        F : Float := 0.0;
6
        I : Integer := 0;
7
      end record
8
        with Dynamic_Predicate =>
9
                Integer (Float_Integer.F) =
10
                  Float_Integer.I,
11
              Put Image
                                => Put_Float_Integer;
12
              ^ Custom version of Put_Image
13
14
      use Ada.Strings.Text_Buffers;
15
16
      procedure Put_Float_Integer
17
         (Buffer : in out Root_Buffer_Type'Class;
18
                          Float Integer);
         Arg
                 :
19
20
   end Custom Numerics;
21
```

¹⁵⁰ http://www.ada-auth.org/standards/22rm/html/RM-4-10.html

Listing 40: custom_numerics.adb

```
package body Custom_Numerics is
1
2
      procedure Put Float Integer
3
        (Buffer : in out Root Buffer Type Class;
4
         Arg
                          Float Integer) is
5
      begin
6
         -- Call Wide Wide Put with customized
7
         -- information
8
         Buffer.Wide_Wide_Put
9
            ("(F : " & Arg.F'Wide_Wide_Image & ", "
10
             & "I : " & Arg.I'Wide_Wide_Image & ")");
11
      end Put_Float_Integer;
12
13
   end Custom Numerics;
14
```

Listing 41: show_put_image.adb

```
with Ada.Text_I0;
                           use Ada.Text IO;
1
2
   with Custom Numerics; use Custom Numerics;
3
4
   procedure Show_Put_Image is
5
      V : Float_Integer;
6
   begin
7
      V := (F => 100.2,
8
             I => 100);
9
      Put_Line ("V =
10
                 & V'Image);
11
   end Show_Put_Image;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Put_Image.Put_Image_Custom_

→Numerics

MD5: 1dbb5fa612b5ca86facc3e93b47977e0
```

Runtime output

V = (F : 1.00200E+02, I : 100)

In the Custom_Numerics package of this example, we specify the Put_Image aspect and indicate the Put_Float_Integer procedure. In that procedure, we display the information of components F and I. Then, in the Show_Put_Image procedure, we use the Image attribute for the V variable and see the information in the exact format we specified. (If you like to see the default version of the Put_Image instead, you may comment out the Put_Image aspect part in the declaration of Float_Integer.)

7.6.3 Relation to the Image attribute

Note that we cannot override the Image attribute directly — there's no Image *aspect* that we could specify. However, as we've just seen, we can do this indirectly by using our own version of the Put_Image procedure for a type T.

The Image attribute of a type T makes use of the procedure indicated in the Put_Image aspect. Let's say we have the following declaration:

```
type T is null record
with Put_Image => Put_Image_T;
```

When we then use the T'Image attribute in our code, the custom Put_Image_T procedure is automatically called. This is a simplified example of how the Image function is implemented:

In other words, the Image attribute basically:

 calls the Put_Image procedure specified in the Put_Image aspect of type T's declaration and passes a buffer;

and

• retrieves the contents of the buffer as a string and returns it.

If the Put_Image aspect of type T isn't specified, the default version is used. (We've seen the default version of various types *in the previous section* (page 338) about the Image attribute.)

7.6.4 Put_Image and derived types

Types that were derived from untagged types (or null extensions) make use of the Put_Image procedure that was specified for their parent type — either a custom procedure indicated in the Put_Image aspect or the default one. Naturally, if a derived type has the Put_Image aspect, the procedure indicated in the aspect is used instead. For example:

Listing 42: untagged_put_image.ads

```
with Ada.Strings.Text_Buffers;
1
2
   package Untagged_Put_Image is
3
4
      use Ada.Strings.Text Buffers;
5
6
      type T is null record
7
        with Put_Image => Put_Image_T;
8
9
      procedure Put Image T
10
         (Buffer : in out Root_Buffer_Type'Class;
11
                          T);
12
         Arg
                 .
13
      type T_Derived_1 is new T;
14
15
      type T Derived 2 is new T
16
        with Put Image => Put Image T Derived 2;
17
18
      procedure Put_Image_T_Derived_2
19
         (Buffer : in out Root Buffer Type'Class;
20
                           T_Derived_2);
         Ara
                 :
21
22
   end Untagged_Put_Image;
23
```

Listing 43: untagged_put_image.adb

```
package body Untagged_Put_Image is
1
2
      procedure Put Image T
3
         (Buffer : in out Root Buffer Type Class;
4
                           T) is
          Arg
5
          pragma Unreferenced (Arg);
6
      begin
7
          Buffer.Wide Wide Put ("Put Image T");
8
      end Put_Image_T;
9
10
      procedure Put_Image_T_Derived_2
11
         (Buffer : in out Root_Buffer_Type'Class;
12
                           T Derived 2) is
         Ara
13
          pragma Unreferenced (Arg);
14
      begin
15
          Buffer.Wide Wide Put
16
            ("Put_Image_T_Derived_2");
17
      end Put_Image_T_Derived_2;
18
19
   end Untagged_Put_Image;
20
```

Listing 44: show_untagged_put_image.adb

```
with Ada.Text_I0;
                              use Ada.Text_I0;
1
2
   with Untagged_Put_Image; use Untagged_Put_Image;
3
4
   procedure Show_Untagged_Put_Image is
5
      Obj_T
                       : T;
6
      Obj_T_Derived_1 : T_Derived_1;
7
      Obj_T_Derived_2 : T_Derived_2;
8
   begin
9
      Put_Line ("T'Image :
10
                 & Obj_T'Image);
11
      Put_Line ("T_Derived_1'Image : "
12
                 & Obj T Derived 1'Image);
13
      Put_Line ("T_Derived_2'Image : "
14
                 & Obj_T_Derived_2'Image);
15
   end Show Untagged Put Image;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Strings.Put_Image.Untagged_Put_Image
MD5: acc0c17d45e6271cb582e65bfc8a2a98
```

Runtime output

T'Image : Put_Image_T
T_Derived_1'Image : Put_Image_T
T_Derived_2'Image : Put_Image_T_Derived_2

In this example, we declare the type T and its derived types T_Derived_1 and T_Derived_2. When running this code, we see that:

- T_Derived_1 makes use of the Put_Image_T procedure from its parent.
 - Note that, if we remove the Put_Image aspect from the declaration of T, the default version of the Put_Image procedure is used for both T and T_Derived_1 types.

• T_Derived_2 makes use of the Put_Image_T_Derived_2 procedure, which was indicated in the Put_Image aspect of that type, instead of its parent's procedure.

7.6.5 Put_Image and tagged types

Types that are derived from a tagged type may also inherit the Put_Image aspect. However, there are a couple of small differences in comparison to untagged types, as we can see in the following example:

Listing 45: tagged_put_image.ads

```
with Ada.Strings.Text_Buffers;
1
2
   package Tagged Put Image is
3
4
      use Ada.Strings.Text_Buffers;
5
6
      type T is tagged record
7
         I : Integer := 0;
8
      end record
9
        with Put_Image => Put_Image_T;
10
11
      procedure Put_Image T
12
         (Buffer : in out Root_Buffer_Type'Class;
13
                           T);
          Ara
                 :
14
15
      type T_Child_1 is new T with record
16
          I1 : Integer;
17
      end record;
18
19
      type T_Child_2 is new T with null record;
20
21
      type T_Child_3 is new T with record
22
          I3 : Integer := 0;
23
      end record
24
        with Put_Image => Put_Image_T_Child_3;
25
26
      procedure Put_Image_T_Child_3
27
         (Buffer : in out Root_Buffer_Type'Class;
28
                           T_Child_3);
29
          Arg
30
   end Tagged_Put_Image;
31
```

Listing 46: tagged_put_image.adb

```
package body Tagged_Put_Image is
1
2
      procedure Put_Image_T
3
         (Buffer : in out Root_Buffer_Type'Class;
4
         Arg
                          T) is
5
         pragma Unreferenced (Arg);
6
7
      begin
         Buffer.Wide_Wide_Put ("Put_Image_T");
8
      end Put_Image_T;
9
10
      procedure Put_Image_T_Child_3
11
         (Buffer : in out Root_Buffer_Type'Class;
12
                           T Child 3) is
         Arg
13
                 .
         pragma Unreferenced (Arg);
14
      begin
15
         Buffer.Wide_Wide_Put
16
```

(continues on next page)

(continued from previous page)

```
17 ("Put_Image_T_Child_3");
18 end Put_Image_T_Child_3;
19
20 end Tagged_Put_Image;
```

Listing 47: show_tagged_put_image.adb

```
with Ada.Text I0;
                         use Ada.Text I0;
1
2
   with Tagged Put Image; use Tagged Put Image;
3
4
   procedure Show_Tagged_Put_Image is
5
      Obj_T
                  : T;
6
      Obj_T_Child 1 : T Child 1;
7
      Obj T Child 2 : T Child 2;
8
      Obj_T_Child_3 : T_Child_3;
9
   begin
10
      Put Line ("T'Image :
11
               & Obj_T'Image);
12
      Put Line ("-----");
13
      Put_Line ("T_Child_1'Image : "
14
               & Obj_T_Child_1'Image);
15
      Put Line ("-----
16
                                     ');
      Put_Line ("T_Child_2'Image : "
17
               & Obj_T_Child_2'Image);
18
      Put Line ("-----");
19
      Put Line ("T Child 3'Image : "
20
               & Obj_T_Child_3'Image);
21
      Put Line ("-----");
22
      Put_Line ("T'Class'Image : "
23
               & T'Class (Obj_T_Child_1)'Image);
24
   end Show_Tagged_Put_Image;
25
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Strings.Put_Image.Tagged_Put_Image
MD5: b19214bbcbc8c0339ead744afffcdd68

Runtime output

```
T'Image : Put_Image_T
T_Child_1'Image :
(Put_Image_T with I1 => 2)
T_Child_2'Image : Put_Image_T
T_Child_3'Image : Put_Image_T_Child_3
T'Class'Image : TAGGED_PUT_IMAGE.T_CHILD_1'
(Put_Image_T with I1 => 2)
```

In this example, we declare the type T and its derived types T_Child_1, T_Child_2 and T_Child_3. When running this code, we see that:

- for both T_Child_1 and T_Child_2, the parent's Put_Image aspect (the Put_Image_T procedure) is called and its information is combined with the information from the type extension;
 - For the T_Child_1 type, the I1 component of the type extension is displayed by calling a default version of the Put_Image procedure for that component — (Put_Image_T with I1 => 0) is displayed.

- For the T_Child_2 type, no additional information is displayed because this type has a null extension.
- for the T_Child_3 type, the Put_Image_T_Child_3 procedure, which was indicated in the Put_Image aspect of the type, is used.

Finally, class-wide types (such as T'Class) include additional information. Here, the tag of the specific derived type is displayed first — in this case, the tag of the T_Child_1 type — and then the actual information for the derived type is displayed.

7.7 Universal text buffer

In the *previous section* (page 345), we've seen that the first parameter of the procedure indicated in the Put_Image aspect has the Root_Buffer_Type'Class type, which is defined in the Ada.Strings.Text_Buffers package. In this section, we talk more about this type and additional procedures associated with this type.

\rm 1 Note

This feature was introduced in Ada 2022.

7.7.1 Overview

We use the Root_Buffer_Type 'Class type to implement a universal text buffer that is used to store and retrieve information about data types. Because this text buffer isn't associated with specific data types, it is universal — in the sense that we can really use it for any data type, regardless of the characteristics of this type.

In theory, we could use Ada's universal text buffer to implement applications that actually process text in some form — for example, when implementing a text editor. However, in general, Ada programmers are only expected to make use of the Root_Buffer_Type'Class type when implementing a procedure for the Put_Image aspect. For this reason, we won't discuss any kind of type derivation — or any other kind of usages of this type — in this section. Instead, we'll just focus on additional subprograms from the Ada.Strings. Text_Buffers package.

1 In the Ada Reference Manual

• Universal Text Buffers¹⁵¹

7.7.2 Additional procedures

In the previous section, we used the Put procedure — and the related Wide_Put and Wide_Wide_Put procedures — from the Ada.Strings.Text_Buffers package. In addition to these procedures, the package also includes:

- the New_Line procedure, which writes a new line marker to the text buffer;
- the Increase_Indent procedure, which increases the indentation in the text buffer; and
- the Decrease_Indent procedure, which decreases the indentation in the text buffer.

The Ada.Strings.Text_Buffers package also includes the Current_Indent function, which retrieves the current indentation counter.

Let's revisit an example from the previous section and use the procedures mentioned above:

¹⁵¹ http://www.ada-auth.org/standards/22rm/html/RM-A-4-12.html

Listing 48: custom_numerics.ads

```
with Ada.Strings.Text_Buffers;
1
2
   package Custom_Numerics is
3
4
      type Float_Integer is record
5
         F : Float;
6
         I : Integer;
7
      end record
8
        with Dynamic Predicate =>
9
                Integer (Float_Integer.F) =
10
                  Float_Integer.I,
11
              Put_Image
                                 => Put_Float_Integer;
12
              ^ Custom version of Put_Image
13
       - -
14
      use Ada.Strings.Text_Buffers;
15
16
      procedure Put Float Integer
17
         (Buffer : in out Root_Buffer_Type'Class;
18
                           Float Integer);
         Arg
                 .
19
20
   end Custom_Numerics;
21
```

Listing 49: custom_numerics.adb

```
package body Custom_Numerics is
1
2
3
      procedure Put_Float_Integer
         (Buffer : in out Root_Buffer_Type'Class;
4
          Arg
                           Float_Integer) is
5
                  :
      begin
6
          Buffer.Wide_Wide_Put ("(");
7
          Buffer.New_Line;
8
9
          Buffer.Increase_Indent;
10
11
          Buffer.Wide_Wide_Put
12
            ("F : "
13
             & Arg.F'Wide_Wide_Image);
14
          Buffer.New_Line;
15
16
          Buffer.Wide_Wide_Put
17
            ("I : "
18
            & Arg.I'Wide_Wide_Image);
19
20
          Buffer.Decrease_Indent;
21
          Buffer.New_Line;
22
23
          Buffer.Wide_Wide_Put (")");
24
       end Put_Float_Integer;
25
26
   end Custom_Numerics;
27
```

Listing 50: show_put_image.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Custom_Numerics; use Custom_Numerics;
procedure Show_Put_Image is
```

(continues on next page)

(continued from previous page)

Code block metadata

Runtime output

```
V = (
F : 1.00200E+02
I : 100
)
```

In the body of the Put_Float_Integer procedure, we're using the New_Line, Increase_Indent and Decrease_Indent procedures to improve the format of the string returned by the Float_Integer'Image attribute. Using these procedures, you can create any kind of output format for your custom type.

CHAPTER EIGHT

NUMERICS

8.1 Numeric Literals

8.1.1 Classification

We've already discussed basic characteristics of numeric literals in the Introduction to Ada course — although we haven't used this terminology there. There are two kinds of numeric literals in Ada: integer literals and real literals. They are distinguished by the absence or presence of a radix point. For example:

Listing 1: real_integer_literals.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Real_Integer_Literals is
3
      Integer_Literal : constant := 365;
4
      Real_Literal : constant := 365.2564;
5
   begin
6
      Put_Line ("Integer Literal: "
7
                & Integer_Literal'Image);
8
      Put_Line ("Real Literal:
9
                & Real_Literal'Image);
10
   end Real_Integer_Literals;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Real_Integer_

→Literals

MD5: balcc348cad054f3ab86c05e051b40fa
```

Runtime output

```
Integer Literal: 365
Real Literal: 3.6525640000000000E+02
```

Another classification takes the use of a base indicator into account. (Remember that, when writing a literal such as 2#1011#, the base is the element before the first # sign.) So here we distinguish between decimal literals and based literals. For example:

Listing 2: decimal_based_literals.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Decimal_Based_Literals is
package F_I0 is new
Ada.Text_I0.Float_I0 (Float);
```

(continues on next page)

8

(continued from previous page)

```
-- DECIMAL LITERALS
9
10
       - -
11
       Dec_Integer : constant := 365;
12
13
       Dec_Real
                  : constant := 365.2564;
14
       Dec_Real_Exp : constant := 0.365_256_4e3;
15
16
17
       - -
          BASED LITERALS
18
19
20
       Based_Integer
                        : constant := 16#16D#;
21
       Based_Integer_Exp : constant := 5#243#e1;
22
23
       Based Real
                          : constant :=
24
        2#1_0110_1101.0100_0001_1010_0011_0111#;
25
       Based Real Exp : constant :=
26
        7#1.031_153_643#e3;
27
   begin
28
       F_IO.Default_Fore := 3;
29
       F_IO.Default_Aft := 4;
30
       F_IO.Default_Exp := 0;
31
32
       Put_Line ("Dec_Integer:
33
                 & Dec_Integer'Image);
34
35
       Put ("Dec_Real:
                                   "):
36
       F_IO.Put (Item => Dec_Real);
37
       New_Line;
38
39
                                 ");
       Put ("Dec_Real_Exp:
40
       F_I0.Put (Item => Dec_Real_Exp);
41
       New_Line;
42
43
                                       н
       Put_Line ("Based_Integer:
44
                 & Based_Integer'Image);
45
       Put_Line ("Based_Integer_Exp:"
46
                 & Based_Integer_Exp'Image);
47
48
       Put ("Based Real:
                                   ");
49
       F_IO.Put (Item => Based_Real);
50
       New_Line;
51
52
       Put ("Based_Real_Exp:
                                  ");
53
       F_IO.Put (Item => Based_Real_Exp);
54
       New_Line;
55
   end Decimal_Based_Literals;
56
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Decimal_Based_

→Literals

MD5: bde8f422c3844826819348d18fb48a33
```

Runtime output

Dec_Integer:	365
<pre>Dec_Real:</pre>	365.2564
<pre>Dec_Real_Exp:</pre>	365.2564

(continues on next page)

(continued from previous page)

Based_Integer:	365
Based_Integer_Exp:	365
Based_Real:	365.2564
<pre>Based_Real_Exp:</pre>	365.2564

Based literals use the base#number# format. Also, they aren't limited to simple integer literals such as 16#16D#. In fact, we can use a radix point or an exponent in based literals, as well as underscores. In addition, we can use any base from 2 up to 16. We discuss these aspects further in the next section.

8.1.2 Features and Flexibility

1 Note

This section was originally written by Franco Gasperoni and published as Gem #7: The Beauty of Numeric Literals in Ada¹⁵².

Ada provides a simple and elegant way of expressing numeric literals. One of those simple, yet powerful aspects is the ability to use underscores to separate groups of digits. 3.14159 26535 89793 23846 26433 83279 50288 41971 69399 37510 For example, is more readable and less error prone to type than 3. 14159265358979323846264338327950288419716939937510. Here's the complete code:

Listing 3: ada numeric literals.adb

```
with Ada.Text IO;
1
2
    procedure Ada Numeric Literals is
3
       Pi : constant :=
4
         3.14159 26535 89793 23846 26433 83279 50288 41971 69399 37510;
5
6
       Pi2 : constant :=
7
         3.14159265358979323846264338327950288419716939937510;
8
9
           : constant := Pi - Pi2;
       7
10
       pragma Assert (Z = 0.0);
11
12
       use Ada.Text I0;
13
    begin
14
       Put Line ("Z = " & Float'Image (Z));
15
    end Ada Numeric Literals;
16
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Pi_Literals
MD5: 8f6516730fa98f08234b159488431aaf

Runtime output

Z = 0.00000E+00

Also, when using based literals, Ada allows any base from 2 to 16. Thus, we can write the decimal number 136 in any one of the following notations:

¹⁵² https://www.adacore.com/gems/ada-gem-7

Listing 4: ada	numeric	literals.adb
----------------	---------	--------------

```
with Ada.Text IO;
1
2
   procedure Ada_Numeric_Literals is
3
       Bin_136 : constant := 2#1000_1000#;
4
       Oct_136 : constant := 8#210#;
5
       Dec_136 : constant := 10#136#;
6
       Hex 136 : constant := 16#88#;
7
       pragma Assert (Bin_136 = 136);
8
       pragma Assert (Oct_136 = 136);
pragma Assert (Dec_136 = 136);
9
10
       pragma Assert (Hex_136 = 136);
11
12
       use Ada.Text_I0;
13
14
   begin
15
       Put Line ("Bin 136 = "
16
                  & Integer'Image (Bin_136));
17
       Put Line ("Oct 136 = "
18
                  & Integer'Image (Oct 136));
19
       Put_Line ("Dec_136 = "
20
                  & Integer'Image (Dec_136));
21
       Put Line ("Hex 136 = "
22
                  & Integer'Image (Hex_136));
23
   end Ada_Numeric_Literals;
24
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Based_Literals
MD5: 0959ec5e4aafcde245c5a15597ac9b7e

Runtime output

Bin_136 = 136 Oct_136 = 136 Dec_136 = 136 Hex_136 = 136

1 In other languages

The rationale behind the method to specify based literals in the C programming language is strange and unintuitive. Here, you have only three possible bases: 8, 10, and 16 (why no base 2?). Furthermore, requiring that numbers in base 8 be preceded by a zero feels like a bad joke on us programmers. For example, what values do 0210 and 210 represent in C?

When dealing with microcontrollers, we might encounter I/O devices that are memory mapped. Here, we have the ability to write:

Lights_On : constant := 2#1000_1000#; Lights_Off : constant := 2#0111_0111#;

and have the ability to turn on/off the lights as follows:

Output_Devices := Output_Devices or Lights_On; Output_Devices := Output_Devices and Lights_Off;

Here's the complete example:

```
Listing 5: ada_numeric_literals.adb
```

```
with Ada.Text IO;
1
2
   procedure Ada_Numeric_Literals is
3
      Lights_On : constant := 2#1000_1000#;
4
      Lights Off : constant := 2#0111 0111#;
5
6
      type Byte is mod 256;
7
      Output Devices : Byte := 0;
8
9
      -- for Output Devices'Address
10
           use 16#DEAD BEEF#;
      - -
11
          _____
      - -
12
      -- Memory mapped Output
13
14
      use Ada.Text_I0;
15
   begin
16
      Output Devices := Output Devices or
17
                           Lights_On;
18
19
      Put_Line ("Output_Devices (lights on ) = "
20
                & Byte'Image (Output_Devices));
21
22
      Output_Devices := Output_Devices and
23
                           Lights_Off;
24
25
      Put Line ("Output Devices (lights off) = "
26
                & Byte'Image (Output Devices));
27
   end Ada_Numeric_Literals;
28
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Literal_Lights
MD5: c3e72b25366d8d815a1f425f2695ad0b

Runtime output

Output_Devices (lights on) = 136 Output_Devices (lights off) = 0

Of course, we can also use *records with representation clauses* (page 99) to do the above, which is even more elegant.

The notion of base in Ada allows for exponents, which is particularly pleasant. For instance, we can write:

Listing 6:	literal	binaries.ads

```
package Literal Binaries is
1
        Kilobyte : constant := 2#1#e+10;
Megabyte : constant := 2#1#e+20;
Gigabyte : constant := 2#1#e+30;
Terabyte : constant := 2#1#e+40;
2
3
4
5
         Petabyte : constant := 2#1#e+50;
6
         Exabyte : constant := 2#1#e+60;
7
         Zettabyte : constant := 2#1#e+70;
8
         Yottabyte : constant := 2#1#e+80;
9
   end Literal Binaries;
10
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Literal_Binary
MD5: 98d971e0f170db570069f8868e442d6d

In based literals, the exponent — like the base — uses the regular decimal notation and specifies the power of the base that the based literal should be multiplied with to obtain the final value. For instance $2\#1\#e+10 = 1 \times 2^{10} = 1_{024}$ (in base 10), whereas $16\#F\#e+2 = 15 \times 16^2 = 15 \times 256 = 3$ 840 (in base 10).

Based numbers apply equally well to real literals. We can, for instance, write:

One_Third : constant := 3#0.1#; -- same as 1.0/3

Whether we write 3#0.1# or 1.0 / 3, or even 3#1.0#e-1, Ada allows us to specify exactly rational numbers for which decimal literals cannot be written.

The last nice feature is that Ada has an open-ended set of integer and real types. As a result, numeric literals in Ada do not carry with them their type as, for example, in C. The actual type of the literal is determined from the context. This is particularly helpful in avoiding overflows, underflows, and loss of precision.

1 In other languages

In C, a source of confusion can be the distinction between 321 and 321. Although both look similar, they're actually very different from each other.

And this is not all: all constant computations done at compile time are done in infinite precision, be they integer or real. This allows us to write constants with whatever size and precision without having to worry about overflow or underflow. We can for instance write:

Zero : constant := 1.0 - 3.0 * One_Third;

and be guaranteed that constant Zero has indeed value zero. This is very different from writing:

where Zero_Approx is really 1.0e-29 — and that will show up in your numerical computations. The above is quite handy when we want to write fractions without any loss of precision. Here's the complete code:

Listing 7: ada_numeric_literals.adb

```
with Ada.Text I0;
1
2
   procedure Ada_Numeric_Literals is
3
     One Third : constant := 3#1.0#e-1;
4
     -- same as 1.0/3.0
5
6
     Zero : constant := 1.0 - 3.0 * One_Third;
7
     pragma Assert (Zero = 0.0);
8
9
     One Third Approx : constant :=
10
       11
     Zero Approx
                  : constant :=
12
       1.0 - 3.0 * One Third Approx;
13
```

(continues on next page)

(continued from previous page)

```
14
       use Ada.Text_I0;
15
16
   begin
17
                                = -^{0}
       Put_Line ("Zero
18
                  & Float'Image (Zero));
19
       Put_Line ("Zero_Approx =
20
                  & Float'Image (Zero_Approx));
21
   end Ada_Numeric_Literals;
22
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Literals
MD5: ee604245b34e8cb878a8ebdb21cd564e

Runtime output

Zero = 0.00000E+00 Zero_Approx = 1.00000E-29

Along these same lines, we can write:

Listing 8: ada_numeric_literals.adb

```
with Ada.Text_I0;
1
2
   with Literal Binaries; use Literal Binaries;
3
4
   procedure Ada Numeric Literals is
5
6
      Big Sum : constant := 1
7
                                         +
8
                              Kilobyte +
9
                              Megabyte +
                              Gigabyte +
10
                              Terabyte +
11
                              Petabyte +
12
                              Exabvte
                                         +
13
                              Zettabyte;
14
15
      Result : constant := (Yottabyte - 1) /
16
                              (Kilobyte - 1);
17
18
      Nil
             : constant := Result - Big Sum;
19
      pragma Assert (Nil = 0);
20
21
      use Ada.Text_I0;
22
23
   begin
24
                               = "
      Put Line ("Nil
25
                 & Integer'Image (Nil));
26
   end Ada Numeric Literals;
27
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Numeric_Literals.Literal_Binary
MD5: 7bda6442e68271d12bdb827b63f0d702

Runtime output

```
Nil =
```

and be guaranteed that Nil is equal to zero.

0

8.2 Universal Numeric Types

Previously, we introduced the concept of *universal types* (page 28). Three of them are numeric types: universal real, universal integer and universal fixed types. In this section, we discuss these universal numeric types in more detail.

8.2.1 Universal Real and Integer

Universal real and integer types are mainly used in the declaration of *named numbers* (page 13):

Listing 9: show_universal_real_integer.ads

```
package Show Universal Real Integer is
1
2
      Pi : constant := 3.1415926535;
3
                          ~~~~~
4
       - -
       - -
                     universal real type
5
6
      N : constant := 10;
7
8
       - -
                 universal integer type
9
       - -
10
  end Show Universal Real Integer;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.

⊲Universal_Real_Integer

MD5: 3cfa52af66185c693ede07f3b0e689e6
```

The type of a named number is implied by the type of the *numeric literal* (page 355) and the type of any named numbers that we use in the *static expression* (page 362). (We discuss static expressions next.) In this specific example, we declare Pi using a real literal, which implies that it's a named number of universal real type. Likewise, N is of universal integer type because we use an integer literal in its declaration.

In the Ada Reference Manual

• 3.3.2 Number Declarations¹⁵³

Static expressions

As we've just seen, we can use an expression in the declaration of a named number. This expression is static, as it's always evaluated at compile time. Therefore, we must use the keyword **constant** in the declaration of named numbers.

If all components of the static expression are of universal integer type, then the named number is of universal integer type. Otherwise, the static expression is of universal real type. For example, if the first element of a static expression is of universal integer type, but we have a constant of universal real type in the same expression, then the type of the whole static expression is universal real:

¹⁵³ http://www.ada-auth.org/standards/22rm/html/RM-3-3-2.html

Listing 10: static_expressions.ads

```
package Static Expressions is
1
2
      Two Pi : constant := 2 * 3.1415926535;
3
4
       - -
       - -
                        universal integer type
5
       - -
6
                                  ~~~~~~
       - -
7
                           universal real type
8
       - -
9
       - -
               => result: universal real type
       - -
10
11
   end Static Expressions;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.Static_

⇔Expressions

MD5: 3429db9e1a7c4d4fe7d94e82159c3cb8
```

In this example, the static expression is of universal real type because of the real literal (3.1415926535) — even though we have the universal integer 2 in the expression.

Likewise, if we use a constant of universal real type in the static expression, the result is of universal real type:

Listing 11:	static	_expressions.ads
-------------	--------	------------------

```
package Static Expressions is
1
2
       Pi
              : constant := 3.1415926535;
3
                                     ~~~~
4
                          universal real type
5
       - -
6
       Two_Pi : constant := 2 * Pi;
7
8
       - -
       - -
                         universal integer type
9
       - -
10
                                   ~~
       - -
11
       - -
                            universal real type
12
       - -
13
                => result: universal real type
       - -
14
15
   end Static_Expressions;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.Static_

→Expressions

MD5: 599494102c4e5e5979a6e0071412da78
```

In this example, the result of the static expression is of universal real type because of we're using the named number Pi, which is of universal real type.

Complexity of static expressions

The operations that we use in static expressions may be arbitrarily complex. For example:

LISUNG 12: SLAUC EXPRESSIONS.au	Listing	tatic expressions.ac	static e
---------------------------------	---------	----------------------	----------

```
package Static Expressions is
1
2
       C1 : constant := 300 453.5;
3
       C2 : constant := 455_233.5 * C1;
4
       C3 : constant := 872_922.5 * C2;
5
       C4 : constant := 155_277.5 * C1 + C2 / C3;
6
       C5 : constant := 2.0 * C1 +
7
                          3.0 * (C2 / (C4 * C3)) +
8
                          4.0 * (C1 / (C2 * C2)) +
5.0 * (C3 / (C1 * C1));
9
10
11
   end Static_Expressions;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.Static_

→Expressions

MD5: ebdd5b1c64ad1944931a962756e72291
```

As we can see in this example, we may create a chain of dependencies, where the result of a static expression depends on the result of previously evaluated static expressions. For instance, C5 depends on the evaluation of C1, C2, C3, C4.

Accuracy of static expressions

The accuracy and range of numeric literals used in static expressions may be arbitrarily high as well:

Listing 13:	static	expressions.ads
-------------	--------	-----------------

```
package Static_Expressions is
1
2
      Pi : constant :=
3
         3.14159 26535 89793 23846 26433 83279 50288;
4
5
      Seed : constant :=
6
         143_574_786_272_784_656_928_283_872_972_764;
7
8
      Super Seed : constant :=
9
         Seed * Seed * Seed * Seed * Seed * Seed;
10
11
   end Static Expressions;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.Static_

⇔Expressions

MD5: 777574a29ffa6da8bffb4287dee45be8
```

In this example, Super_Seed has a value that is above the typical range of integer constants. This might become challenging when using such named numbers in actual computations, as we *discuss soon* (page 368).

Another example is when the result of the expression is a repeating decimal¹⁵⁴:

¹⁵⁴ https://en.wikipedia.org/wiki/Repeating_decimal

Listing 14: repeating_decimals.ads

```
package Repeating_Decimals is
    One_Over_Three : constant :=
        1.0 / 3.0;
    end Repeating Decimals;
```

Listing 15: show_repeating_decimals.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Repeating Decimals;
3
   use Repeating_Decimals;
4
5
   procedure Show_Repeating_Decimals is
6
      F_1_3
               : constant Float
7
                                             :=
                    One_Over_Three;
8
9
      LF_1_3
                : constant Long_Float
                                             :=
                    One_Over_Three;
10
      LLF_1_3 : constant Long_Long_Float :=
11
                    One_Over_Three;
12
   begin
13
      Put_Line (F_1_3'Image);
14
      Put_Line (LF_1_3'Image);
15
      Put_Line (LLF_1_3'Image);
16
   end Show_Repeating_Decimals;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.

⇔Repeating_Decimal

MD5: 4fc38ef6482e403d655b4662d4199abb
```

Runtime output

```
3.33333E-01
3.33333333333333E-01
3.3333333333333333333E-01
```

In this example, as expected, we see that the accuracy of the value we display increases if we use a type with higher precision. This wouldn't be possible if we had used a floating-point type with limited precision for the 0ne 0ver Three constant:

```
Listing 16: repeating_decimals.ads
```

```
package Repeating Decimals is
1
2
      One_Over_Three : constant Long_Float :=
3
         1.0 / 3.0;
4
                                 ~~~~~~
5
                  using Long Float instead of
      - -
6
                      universal real type
7
8
   end Repeating Decimals;
9
```

Listing 17: show_repeating_decimals.adb

```
with Ada.Text_I0; use Ada.Text_I0;
```

(continues on next page)

(continued from previous page)

```
with Repeating_Decimals;
3
   use Repeating_Decimals;
4
5
   procedure Show_Repeating_Decimals is
6
      F_1_3 : constant Float
                                            :=
7
                    Float (One_Over_Three);
8
      LF_1_3 : constant Long_Float
9
                                            :=
                    Long_Float (One_Over_Three);
10
      LLF_1_3 : constant Long_Long_Float :=
11
                    Long_Long_Float (One_Over_Three);
12
   begin
13
      Put_Line (F_1_3'Image);
14
      Put_Line (LF_1_3'Image);
15
      Put_Line (LLF_1_3'Image);
16
   end Show_Repeating_Decimals;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.

←Repeating_Decimal

MD5: d0fa105d679cc246e2e8baf37cbe48c4
```

Runtime output

```
3.33333E-01
3.333333333333333E-01
3.333333333333333333315E-01
```

Because we're using the Long_Float type for the One_Over_Three constant instead of the universal real type, the accuracy doesn't increase when we use the Long_Long_Float type — as we see in the value of the LLF_1_3 constant — even though this type has a higher precision.

1 For further reading...

When using *big numbers* (page 408), you could simply assign the named number One_Over_Three to a big real:

Listing 18: repeating_decimals.ads

```
package Repeating_Decimals is
```

```
One_Over_Three : constant :=
    1.0 / 3.0;
end Repeating Decimals;
```

1 2

3

4 5

6

```
Listing 19: show_repeating_decimals.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
```

```
with Ada.Numerics.Big_Numbers.Big_Reals;
use Ada.Numerics.Big_Numbers.Big_Reals;
```

```
with Repeating_Decimals;
use Repeating Decimals;
```

procedure Show_Repeating_Decimals is
 BR_1_3 : constant Big_Real := One_Over_Three;
begin

end Show_Repeating_Decimals;

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.

→Repeating_Decimal

MD5: 4f1981b785baa35704c85e7e688c8ce4
```

Runtime output

Another approach is to use the division operation directly:

Listing 20: show_repeating_decimals.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Ada.Numerics.Big_Numbers.Big_Reals;
use Ada.Numerics.Big_Numbers.Big_Reals;
with Repeating_Decimals;
use Repeating_Decimals;
procedure Show Repeating Decimals is
   BR 1 3
           : constant Big Real := 1 / 3;
begin
   Put Line ("BR: "
             & To_String (Arg => BR_1_3,
                          Fore => 2,
                          Aft => 31,
                          Exp
                               => ());
end Show_Repeating_Decimals;
```

```
2
 3
 4
 5
 6
 7
 8
9
10
11
12
13
14
15
16
17
```

1

1 2

3

4 5

6

7 8

9

10

11

12

13

14

15

16

17

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types. ⊲Repeating_Decimal MD5: 5fc195f9fbab3b1ec74c507780a44ec8

Runtime output

We talk more about big real and quotients (page 422) later on.

Conversion of universal real and integer

Although a named number exists as an numeric representation form in Ada, the value it represents cannot be used directly at runtime — even if we *just* display the value of the constant at runtime, for example. In fact, a conversion to a non-universal type is required in order to use the named number anywhere else other than a static expression:

```
Listing 21: static expressions.ads
```

```
package Static_Expressions is
1
2
      Pi : constant :=
3
         3.14159_26535_89793_23846_26433_83279_50288;
4
5
      Seed : constant :=
6
         143_574_786_272_784_656_928_283_872_972_764;
7
8
      Super Seed : constant :=
9
          Seed * Seed * Seed * Seed * Seed * Seed;
10
11
   end Static_Expressions;
12
```

Listing 22: show_static_expressions.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Static Expressions;
3
   use Static_Expressions;
4
5
   procedure Show_Static_Expressions is
6
   begin
7
      Put Line (Pi'Image);
8
       -- Same as:
-- Put_Line (Float (Pi)'Image);
9
10
11
      Put_Line (Seed'Image);
12
      -- Same as:
13
       -- Put Line (
14
             Long Long Long Integer (Seed)'Image);
       - -
15
   end Show Static Expressions;
16
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types. ⊶Conversion_To_Non_Universal_Types MD5: e50641737f970b935e853ac249dd83d8

Runtime output

```
3.14159265358979324E+00
143574786272784656928283872972764
```

As we see in this example, the named number Pi is converted to **Float** before being used as an actual parameter in the call to Put_Line. Similarly, Seed is converted to Long_Long_Long_Integer.

When we use the Image attribute, the compiler automatically selects a numeric type which has a suitable range for the named number. In the example above, we wouldn't be able to represent the value of Seed with **Integer**, so the compiler selected Long_Long_Long_Integer. Of course, we could have also specified the type by using explicit *type conversions* (page 45) or a *qualified expressions* (page 64):

Listing 23: show_static_expressions.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Static_Expressions;
3
   use Static_Expressions;
4
5
   procedure Show Static Expressions is
6
   begin
7
      Put_Line (Long_Long_Float (Pi)'Image);
8
      Put Line (Long Long Float'(Pi)'Image);
9
   end Show_Static_Expressions;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types.

⊶Conversion_To_Non_Universal_Types

MD5: 18bcc3bffd51ebc1bc98976ed1597f01
```

Runtime output

```
3.14159265358979324E+00
3.14159265358979324E+00
```

Now, we're explicitly converting to Long_Long_Float in the first call to Put_Line and using a qualified expression in the second call to Put_Line.

A conversion is also performed when we use a named number in an object declaration:

```
Listing 24: show_static_expressions.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Static Expressions;
3
   use Static_Expressions;
4
5
   procedure Show Static Expressions is
6
      Two Pi : constant Float := 2.0 * Pi;
7
       -- Same as:
8
      -- Two Pi: constant Float :=
9
      - -
                     2.0 * Float (Pi);
10
11
      Two Pi More Precise :
12
        constant Long_Long_Float := 2.0 * Pi;
13
       -- Same as:
14
          Two Pi More Precise :
       - -
15
            constant Long Long Float :=
       - -
16
               2.0 * Long Long Float (Pi);
       - -
17
   begin
18
      Put_Line (Two_Pi'Image);
19
      Put Line (Two Pi More Precise'Image);
20
   end Show_Static_Expressions;
21
```

Code block metadata

Runtime output

6.28319E+00

6.28318530717958648E+00

In this example, Pi is converted to **Float** in the declaration of Two_Pi because we use the **Float** type in its declaration. Likewise, Pi is converted to **Long_Long_Float** in the declaration of Two_Pi_More_Precise because we use the **Long_Long_Float** type in its declaration. (Actually, the same conversion is performed for each instance of the real literal 2.0 in this example.)

Note that the range of the type we select might not be suitable for the named number we want to use. For example:

Listing 25: show_static_expressions.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Static Expressions;
3
   use Static_Expressions;
4
5
   procedure Show Static Expressions is
6
      Initial_Seed : constant
7
        Long_Long_Integer :=
8
          Super Seed;
۵
   begin
10
      Put Line (Initial Seed'Image);
11
  end Show Static Expressions;
12
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types. ⊶Conversion_To_Non_Universal_Types MD5: 2f8e26fbcd0b5defd94ffef570c0f087

Build output

In this example, we get a compilation error because the range of the Long_Long_Long_Integer type isn't enough to store the value of the Super_Seed.

For further reading...

To circumvent the compilation error in the code example we've just seen, the best alternative to use *big numbers* (page 408) — we discuss this topic later on in this chapter:

```
Listing 26: show static expressions.adb
   with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Ada.Numerics.Big_Numbers.Big_Integers;
3
   use Ada.Numerics.Big_Numbers.Big_Integers;
4
5
   with Static Expressions;
6
   use Static_Expressions;
7
8
   procedure Show_Static_Expressions is
9
10
      Initial_Seed : constant
        Big_Integer :=
11
          Super_Seed;
12
```

13 **begin** 14 Pu

- Put_Line (Initial_Seed'Image);
 end Show Static Expressions;
- 15

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Numeric_Types. ⊲Conversion_To_Non_Universal_Types MD5: bf1511f1b8bf3965baa86b953c56c406

Runtime output

By changing the type from Long_Long_Long_Integer to Big_Integer, we get rid of the compilation error. (The value of Super_Seed — stored in Initial_Seed — is displayed at runtime.)

8.2.2 Universal Fixed

For fixed-point types, we also have a corresponding universal type. However, in contrast to the universal real and integer types, universal fixed types aren't an abstraction used in static expressions, but rather a concept that permeates actual fixed-point types. In fact, for *fixed-point types* (page 399), some operations are accomplished via universal fixed types — for example, the conversion between fixed-point types and the multiplication and division operations.

Let's start by analyzing how floating-point and integer types associate their operations to the specific type of an object. For example, if we have an object A of type **Float** in a multiplication, we cannot just write A * B if we want to multiply A by an object B of another floating-point type — if B is of type **Long_Float**, for example, writing A * B triggers a compilation error. (Otherwise, which precision should be used for the result?) Therefore, we have to convert one of the objects to have matching types:

Listing 27: show_float_multiplication_mismatch.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
  procedure Show_Float_Multiplication_Mismatch is
з
      F : Float
                   := 0.25;
4
     LF : Long Float := 0.50;
5
  begin
6
      F := F * LF;
7
     Put Line ("F = " & F'Image);
8
  end Show Float Multiplication Mismatch;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Types.Float_

→Multiplication

MD5: 88ce3a0f29e2bd31ddfc491557d7f0e3
```

Build output

```
show_float_multiplication_mismatch.adb:7:11: error: invalid operand types for_
operator "*"
show_float_multiplication_mismatch.adb:7:11: error: left operand has type
o"Standard.Float"
show_float_multiplication_mismatch.adb:7:11: error: right operand has type
o"Standard.Long_Float"
gprbuild: *** compilation phase failed
```

This code example fails to compile because of the F * LF operation. (We could correct the code by writing F * **Float** (LF), for example.)

In contrast, for fixed-point types, we can mix objects of different types in a multiplication or division. (In this case, mixing is allowed for the convenience of the programmer.) For example:

Listing 28: normalized_fixed_point_types.ads

```
package Normalized Fixed Point Types is
1
2
      type TQ31 is
3
        delta 2.0 ** (-31)
4
        range -1.0 .. 1.0 - 2.0 ** (-31);
5
6
      type TQ15 is
7
        delta 2.0 ** (-15)
8
         range -1.0 .. 1.0 - 2.0 ** (-15);
9
10
   end Normalized Fixed Point Types;
11
```

Listing 29: show_fixed_multiplication.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Normalized Fixed Point Types;
3
   use Normalized Fixed Point Types;
4
5
   procedure Show Fixed Multiplication is
6
      A : TQ15 := 0.25;
7
      B : TQ31 := 0.50;
8
   begin
9
      A := A * B;
10
      Put Line ("A = " & A'Image);
11
  end Show Fixed Multiplication;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Types.Fixed_Point_

⊶Multiplication

MD5: a4cefdc29a562fbec30b6864b6ec2602
```

Runtime output

```
A = 0.12500
```

In this example, the A * B is accepted by the compiler, even though A and B have different types. This is only possible because the multiplication operation of fixed-point types makes use of the universal fixed type. In other words, the multiplication operation in this code example doesn't operate directly on the fixed-point type TQ31. Instead, it converts A and B to the universal fixed type, performs the operation using this type, and converts back to the original type — TQ15 in this case.

In addition to the multiplication operation, other operations such as the conversion between fixed-point types and the division operations make use of universal fixed types:

Listing 30: custom_decimal_types.ads

```
package Custom_Decimal_Types is
type T3_D3 is delta 10.0 ** (-3) digits 3;
type T3_D6 is delta 10.0 ** (-3) digits 6;
```

(continues on next page)

(continued from previous page)

```
type T6_D6 is delta 10.0 ** (-6) digits 6;
end Custom_Decimal_Types;
```

5

7

Listing 31: show universal fixed.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Custom_Decimal_Types;
3
   use Custom_Decimal_Types;
4
5
   procedure Show_Universal_Fixed is
6
     Val T3 D3 : T3 D3;
7
     Val T3 D6 : T3 D6;
8
     Val T6 D6 : T6 D6;
9
   begin
10
      Val T3 D3 := 0.65;
11
12
      Val_T3_D6 := T3_D6 (Val_T3_D3);
13
14
               type conversion using
15
       - -
                universal fixed type
16
       - -
17
      Val_T6_D6 := T6_D6 (Val_T3_D6);
18
19
       - -
               type conversion using
20
                universal fixed type
21
22
      Put_Line ("Val_T3_D3 = "
23
                 & Val_T3_D3'Image);
24
      Put_Line ("Val_T3_D6 = "
25
                 & Val_T3_D6'Image);
26
      Put_Line ("Val_T6_D6 = "
27
                 & Val_T3_D6'Image);
28
      Put_Line ("-----");
29
30
      Val_T3_D6 := Val_T6_D6 * 2.0;
31
32
             using universal fixed type for
       - -
33
               the multiplication operation
34
      Put_Line ("Val_T3_D6 = "
35
                 & Val_T3_D6'Image);
36
37
      Val_T3_D6 := Val_T6_D6 / Val_T3_D3;
38
39
               different fixed-point types:
       - -
40
            using universal fixed type for
       - -
41
                     the division operation
42
      Put Line ("Val T3 D6 = "
43
                 & Val_T3_D6'Image);
44
45
   end Show Universal Fixed;
46
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Types.Universal_Fixed MD5: 1e253d8a39576f817b2130aa35929d96

Runtime output

Val_T3_D3 Val_T3_D0 Val_T6_D0	5 =	0.650 0.650 0.650
Val_T3_D0 Val_T3_D0		1.300 1.000

In this example, the conversion from the fixed-point type T3_D3 to the T3_D6 and T6_D6 types is performed via universal fixed types.

Similarly, the multiplication operation Val_T6_D6 * 2.0 uses universal fixed types. Here, we're actually multiplying a variable of type T6_D6 by two and assigning it to a variable of type Val_T3_D6. Although these variable have different fixed-point types, no explicit conversion (e.g.: Val_T3_D6 := T3_D6 (Val_T6_D6 * 2.0);) is required in this case because the result of the operation is of universal fixed type, so that it can be assigned to a variable of any fixed-point type.

Finally, in the Val_T3_D6 := Val_T6_D6 / Val_T3_D3 statement, we're using three fixedpoint types: we're dividing a variable of type T6_D6 by a variable of type T3_D3, and assigning it to a variable of type T3_D6. All these operations are only possible without explicit type conversions because the underlying types for the fixed-point division operation are universal fixed types.

For further reading...

It's possible to implement custom * and / operators for fixed-point types. However, those operators do **not** override the corresponding operators for universal fixed-point types. For example:

Listing 32: normalized_fixed_point_types.ads

```
package Normalized Fixed Point Types is
1
2
       type TQ63 is
3
         delta 2.0 ** (-63)
4
         range -1.0 .. 1.0 - 2.0 ** (-63);
5
6
       type T031 is
7
         delta 2.0 ** (-31)
8
         range -1.0 .. 1.0 - 2.0 ** (-31);
9
10
      overriding
11
12
           "+" operator is overriding!
       - -
13
       function "+" (L, R : TQ31)
14
                      return T031:
15
16
      not overriding
17
          ~~~~~
18
       -- "*" operator is NOT overriding!
19
       function "*" (L, R : TQ31)
20
                      return T031;
21
22
       type TQ15 is
23
         delta 2.0 ** (-15)
24
         range -1.0 .. 1.0 - 2.0 ** (-15);
25
26
   end Normalized Fixed Point Types;
27
```

```
Listing 33: normalized fixed point types.adb
   with Ada.Text IO; use Ada.Text IO;
1
2
   package body Normalized_Fixed_Point_Types is
3
4
      function "+" (L, R : TQ31)
5
                     return T031 is
6
      begin
7
          Put Line
8
           ("=> Overriding '+'");
9
          return TQ31 (TQ63 (L) + TQ63 (R));
10
      end "+";
11
12
      function "*" (L, R : TQ31)
13
                     return T031 is
14
      begin
15
          Put_Line
16
            ("=> Custom "
17
             & "non-overriding '*'");
18
          return TQ31 (TQ63 (L) * TQ63 (R));
19
20
      end "*";
21
   end Normalized Fixed Point Types;
22
                          Listing 34: show fixed multiplication.adb
   with Ada.Text IO; use Ada.Text IO;
1
2
   with Normalized Fixed Point Types;
3
   use Normalized_Fixed_Point_Types;
4
5
   procedure Show_Fixed_Multiplication is
6
      Q31_A : TQ31 := 0.25;
7
      Q31_B : TQ31 := 0.50;
8
      Q15_A : TQ15 := 0.25;
9
      Q15_B : TQ15 := 0.50;
10
   begin
11
      Q31 A := Q31 A * Q31 B;
12
      Put Line ("Q31 A = " & Q31 A'Image);
13
14
      Q15 A := Q15 A * Q15 B;
15
      Put_Line ("Q15_A = " & Q31_A'Image);
16
17
      Q15_A := TQ15 (Q31_A) * Q15_B;
18
19
       - -
      -- A conversion is required because of
20
      -- the multiplication operator of
21
      -- T015.
22
      Put Line ("Q31 A = " & Q31 A'Image);
23
   end Show Fixed Multiplication;
24
   Code block metadata
   Project: Courses.Advanced_Ada.Data_Types.Numerics.Universal_Types.Fixed_Point_
    →Custom_Multiplication
   MD5: 954ada297ac676ab1f11447083d87882
   Runtime output
   => Custom non-overriding '*'
   Q31_A = 0.1250000000
   Q15_A = 0.1250000000
```

Q31 A = 0.1250000000

In this example, we're declaring a custom multiplication operator for the TQ31 type. As we can see in the declaration, we specify that it's **not overriding** the * operator. (Removing the **not** keyword triggers a compilation error.) In contrast, for the + operator, we're indeed overriding the default + operator of the TQ31 type in the Nor-malized_Fixed_Point_Types because the addition operator is associated with its corresponding fixed-point type, not with the universal fixed-point type. In the Q31_A := Q31_A * Q31_B statement, we see at runtime (through the "=> Custom non-overriding '*'' message) that the custom multiplication is being used.

However, because of this custom * operator, we cannot mix objects of this type with objects of other fixed-point types in multiplication or division operations. Therefore, for a statement such as Q15_A := Q31_A * Q15_B, we have to convert Q31_A to the TQ15 type before multiplying it by Q15_B.

In the Ada Reference Manual

• 4.5.5 Multiplying Operators¹⁵⁵

8.3 Base types

You might remember our discussion on *root types* (page 29) and the corresponding numeric root types.

Ada also has the concept of base types, which *sounds* similar to the concept of the root type. However, the focus of each one is different: while the the root type refers to the derivation tree of a type, the base type refers to the constraints of a type.

In fact, the base type denotes the unconstrained underlying hardware representation selected for a given numeric type. For example, if we were making use of a constrained type T, the compiler would select a type based on the hardware characteristics that has sufficient precision to represent T on the target platform. Of course, that type — the base type — would necessarily be unconstrained.

Let's discuss the **Integer** type as an example. The Ada standard specifies that the minimum range of the **Integer** type is $-2^{**}15 + 1 \dots 2^{**}15 - 1$. In modern 64-bit systems — where wider types such as **Long_Integer** are defined — the range is at least $-2^{**}31 + 1 \dots 2^{**}31 - 1$. Therefore, we could think of the **Integer** type as having the following declaration:

```
type Integer is
  range -2 ** 31 .. 2 ** 31 - 1;
```

However, even though **Integer** is a predefined Ada type, it's actually a subtype of an anonymous type. That anonymous "type" is the hardware's representation for the numeric type as chosen by the compiler based on the requested range (for the signed integer types) or digits of precision (for floating-point types). In other words, these types are actually subtypes of something that does not have a specific name in Ada, and that is not constrained.

In effect,

```
type Integer is
  range -2 ** 31 .. 2 ** 31 - 1;
```

is really as if we said this:

¹⁵⁵ http://www.ada-auth.org/standards/22rm/html/RM-4-5-5.html

```
subtype Integer is
Some_Hardware_Type_With_Sufficient_Range
range -2 ** 31 .. 2 ** 31 - 1;
```

Since the Some_Hardware_Type_With_Sufficient_Range type is anonymous and we therefore cannot refer to it in the code, we just say that **Integer** is a type rather than a subtype.

Let's focus on signed integers — as the other numerics work the same way. When we declare a signed integer type, we have to specify the required range, statically. If the compiler cannot find a hardware-defined or supported signed integer type with at least the range requested, the compilation is rejected. For example, in current architectures, the code below most likely won't compile:

Listing 35: int_def.ads

```
1 package Int_Def is
2
3 type Too_Big_To_Fail is
4 range -2 ** 255 .. 2 ** 255 - 1;
5
6 end Int Def;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Base_Type.Very_Big_Range
MD5: 60c2b9219a0cf080997707a99a0b162b

Build output

int_def.ads:4:07: error: integer type definition bounds out of range
gprbuild: *** compilation phase failed

Otherwise, the compiler maps the named Ada type to the hardware "type", presumably choosing the smallest one that supports the requested range. (That's why the range has to be static in the source code, unlike for explicit subtypes.)

8.3.1 Base

The Base attribute gives us the unconstrained underlying hardware representation selected for a given numeric type. As an example, let's say we declared a subtype of the **Integer** type named One_To_Ten:

Listing 36: my_integers.ads

```
package My_Integers is
subtype One_To_Ten is Integer
range 1 .. 10;
end My_Integers;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Base_Type.Base_Attr MD5: e3f8310ed742e61a65728fecb6caa557

If we then use the Base attribute — by writing One_To_Ten'Base —, we're actually referring to the unconstrained underlying hardware representation selected for One_To_Ten. As One_To_Ten is a subtype of the **Integer** type, this also means that One_To_Ten'Base is equivalent to **Integer**'Base, i.e. they refer to the same base type. (This base type is the

underlying hardware type representing the **Integer** type — but is not the **Integer** type itself.)

The following example shows how the Base attribute affects the bounds of a variable:

```
Listing 37: show_base.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
   with My_Integers; use My_Integers;
2
3
   procedure Show Base is
4
      C : constant One_To_Ten := One_To_Ten'Last;
5
   begin
6
      Using Constrained Subtype : declare
7
         V : One To Ten := C;
8
9
      begin
         Put Line
10
            ("Increasing value for One To Ten...");
11
12
         V := One_To_Ten'Succ (V);
13
      exception
14
         when others =>
15
             Put Line ("Exception raised!");
16
      end Using_Constrained_Subtype;
17
18
      Using Base : declare
19
20
         V : One_To_Ten'Base := C;
21
      begin
         Put Line
22
          ("Increasing value for One_To_Ten'Base...");
23
24
         V := One_To_Ten'Succ (V);
25
      exception
26
         when others =>
27
             Put Line ("Exception raised!");
28
      end Using Base;
29
30
      Put_Line ("One_To_Ten'Last: "
31
                 & One_To_Ten'Last'Image);
32
      Put_Line ("One_To_Ten'Base'Last: '
33
                 & One_To_Ten'Base'Last'Image);
34
   end Show Base;
35
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Base_Type.Base_Attr
MD5: ce3e9fb3ff1619e835e9108ae0a787e7
```

Build output

Runtime output

```
Increasing value for One_To_Ten...
Exception raised!
Increasing value for One_To_Ten'Base...
One_To_Ten'Last: 10
One_To_Ten'Base'Last: 2147483647
```

In the first block of the example (Using_Constrained_Subtype), we're asking for the next

value after the last value of a range — in this case, One_To_Ten'Succ (One_To_Ten'Last). As expected, since the last value of the range doesn't have a successor, a constraint exception is raised.

In the Using_Base block, we're declaring a variable V of One_To_Ten'Base subtype. In this case, the next value exists — because the condition One_To_Ten'Last + 1 <= One_To_Ten'Base'Last is true —, so we can use the Succ attribute without having an exception being raised.

In the following example, we adjust the result of additions and subtractions to avoid constraint errors:

Listing 38: my_integers.ads

```
package My_Integers is
1
2
      subtype One To Ten is Integer range 1 .. 10;
3
4
      function Sat Add (V1, V2 : One To Ten'Base)
5
                          return One_To_Ten;
6
7
      function Sat_Sub (V1, V2 : One_To_Ten'Base)
8
                          return One_To_Ten;
9
10
   end My_Integers;
11
```

Listing 39: my_integers.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
   - -
2
   package body My_Integers is
3
4
       function Saturate (V : One To Ten'Base)
5
                            return One_To_Ten is
6
       begin
7
          -- Put_Line ("SATURATE " & V'Image);
8
9
          if V < One_To_Ten'First then</pre>
10
             return One_To_Ten'First;
11
          elsif V > One_To_Ten'Last then
12
             return One_To_Ten'Last;
13
          else
14
             return V;
15
          end if;
16
17
       end Saturate;
18
       function Sat_Add (V1, V2 : One_To_Ten'Base)
19
                          return One_To_Ten is
20
       begin
21
          return Saturate (V1 + V2);
22
       end Sat_Add;
23
24
       function Sat_Sub (V1, V2 : One_To_Ten'Base)
25
                          return One_To_Ten is
26
       begin
27
          return Saturate (V1 - V2);
28
29
       end Sat_Sub;
30
   end My_Integers;
31
```

Listing 40: show_base.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with My_Integers; use My_Integers;
2
3
   procedure Show Base is
4
5
      type Display Saturate Op is (Add, Sub);
6
7
      procedure Display Saturate
8
         (V1, V2 : One_To_Ten;
9
         0p
                : Display_Saturate_Op)
10
11
      is
          Res : One_To_Ten;
12
      beain
13
          case Op is
14
         when Add =>
15
             Res := Sat Add (V1, V2);
16
         when Sub =>
17
             Res := Sat_Sub (V1, V2);
18
          end case;
19
         Put_Line ("SATURATE " & Op'Image
20
                    & " (" & V1'Image
21
                    & ", " & V2'Image
22
                    & ") = " & Res'Image);
23
      end Display_Saturate;
24
25
   begin
26
      Display Saturate (1, 1, Add);
27
      Display_Saturate (10, 8, Add);
28
      Display_Saturate (1, 8, Sub);
29
   end Show_Base;
30
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Base_Type.Base_Attr_Sat MD5: e9b31345c2efc056bdb71824072852d0

Runtime output

SATURATE ADD (1, 1) = 2 SATURATE ADD (10, 8) = 10 SATURATE SUB (1, 8) = 1

In this example, we're using the Base attribute to declare the parameters of the Sat_Add, Sat_Sub and Saturate functions. Note that the parameters of the Display_Saturate procedure are of One_To_Ten type, while the parameters of the Sat_Add, Sat_Sub and Saturate functions are of the (unconstrained) base subtype (One_To_Ten'Base). In those functions, we perform operations using the parameters of unconstrained subtype and adjust the result — in the Saturate function — before returning it as a constrained value of One_To_Ten subtype.

The code in the body of the My_Integers package contains lines that were commented out — to be more precise, a call to Put_Line call in the Saturate function. If you uncomment them, you'll see the value of the input parameter V (of One_To_Ten'Base type) in the runtime output of the program before it's adapted to fit the constraints of the One_To_Ten subtype.

8.4 Attributes of Modular Types

In the Introduction to Ada course, we've seen that Ada has two kinds of integer type: signed¹⁵⁶ and modular¹⁵⁷ types. For example:

Listing 41: num types.ads

```
1 package Num_Types is
2
3 type Signed_Integer is range 1 .. 1_000_000;
4 type Modular is mod 2**32;
5
6 end Num_Types;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Modular_1
MD5: 2dff9fe22c6bbe52f964befccf68debf
```

In this section, we discuss two attributes of modular types: Modulus and Mod. We also discuss operations on modular types.

1 In the Ada Reference Manual

```
    3.5.4 Integer Types<sup>158</sup>
```

8.4.1 Modulus Attribute

The Modulus attribute returns the modulus of the modular type as a universal integer value. Let's get the modulus of the 32-bit Modular type that we've declared in the Num_Types package of the previous example:

Listing 42: show_modular.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Num_Types; use Num_Types;
procedure Show_Modular is
Modulus_Value : constant := Modular'Modulus;
begin
Put_Line (Modulus_Value'Image);
end Show_Modular;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Modular_1
MD5: 336254ebc8c09ee9921633f6919994fe

Runtime output

4294967296

When we run this example, we get 4294967296, which is equal to 2**32.

¹⁵⁶ https://learn.adacore.com/courses/intro-to-ada/chapters/strongly_typed_language.html#intro-ada-integers
¹⁵⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/strongly_typed_language.html#

intro-ada-unsigned-types

¹⁵⁸ http://www.ada-auth.org/standards/22rm/html/RM-3-5-4.html

8.4.2 Mod Attribute

\rm Note

This section was originally written by Robert A. Duff and published as Gem #26: The Mod Attribute¹⁵⁹.

Operations on signed integers can overflow: if the result is outside the base range, Constraint_Error will be raised. In our previous example, we declared the Signed_Integer type:

```
type Signed_Integer is range 1 .. 1_000_000;
```

The base range of Signed_Integer is the range of Signed_Integer 'Base, which is chosen by the compiler, but is likely to be something like $-2^{**}31$... $2^{**}31$ - 1. (Note: we discussed the Base attribute *in this section* (page 377).)

Operations on modular integers use modular (wraparound) arithmetic. For example:

Listing 43: show_modular.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Num_Types;
                      use Num_Types;
3
4
   procedure Show Modular is
5
      X : Modular;
6
   begin
7
      X := 1;
8
      Put Line (X'Image);
9
10
      X := -X;
11
      Put Line (X'Image);
12
   end Show_Modular;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Modular_1
MD5: e9ac61d2e43585f002fe2b79544ef9d7
```

Runtime output

```
1
4294967295
```

Negating X gives -1, which wraps around to $2^{**}32 - 1$, i.e. all-one-bits.

But what about a type conversion from signed to modular? Is that a signed operation (so it should overflow) or is it a modular operation (so it should wrap around)? The answer in Ada is the former — that is, if you try to convert, say, **Integer**'(-1) to Modular, you will get Constraint_Error:

Listing 44: show_modular.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Num_Types; use Num_Types;
procedure Show_Modular is
```

(continues on next page)

¹⁵⁹ https://www.adacore.com/gems/gem-26

(continued from previous page)

```
6 I : Integer := -1;
7 X : Modular := 1;
8 begin
9 X := Modular (I); -- raises Constraint_Error
10 Put_Line (X'Image);
11 end Show_Modular;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Modular_1
MD5: e8e1a1924efcbe770c719c29547bb863

Build output

Runtime output

raised CONSTRAINT_ERROR : show_modular.adb:9 range check failed

To solve this problem, we can use the **Mod** attribute:

Listing 45: show_modular.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Num Types;
                      use Num Types;
3
4
   procedure Show Modular is
5
      I : constant Integer := -1;
6
      X : Modular := 1;
7
   begin
8
      X := Modular'Mod (I):
9
      Put Line (X'Image);
10
  end Show Modular;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Modular_1
MD5: 572a753de946b7578c5f1b6a795ede98

Runtime output

4294967295

The **Mod** attribute will correctly convert from any integer type to a given modular type, using wraparound semantics.

Historically

In older versions of Ada — such as Ada 95 —, the only way to do this conversion is to use Unchecked_Conversion, which is somewhat uncomfortable. Furthermore, if you're trying to convert to a generic formal modular type, how do you know what size of signed integer type to use? Note that Unchecked_Conversion might malfunction if the source and target types are of different sizes.

The Mod attribute was added to Ada 2005 to solve this problem. Also, we can now safely use this attribute in generics. For example: Listing 46: mod attribute.ads generic

```
type Formal Modular is mod <>;
   package Mod_Attribute is
      function F return Formal_Modular;
4
   end Mod_Attribute;
5
                               Listing 47: mod attribute.adb
   package body Mod_Attribute is
1
2
      A_Signed_Integer : Integer := -1;
3
4
      function F return Formal_Modular is
5
      begin
6
         return Formal_Modular'Mod
                   (A_Signed_Integer);
8
      end F;
9
10
   end Mod Attribute;
11
   Code block metadata
```

Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Mod_Attribute MD5: b2f227b8d4f14cd36508bf33c403f751

In this example, F will return the all-ones bit pattern, for whatever modular type is passed to Formal Modular.

8.4.3 Operations on modular types

Modular types are particularly useful for bit manipulation. For example, we can use the and, or, xor and not operators for modular types.

Also, we can perform bit-shifting by multiplying or dividing a modular object with a power of two. For example, if M is a variable of modular type, then M := M * 2 ** 3; shifts the bits to the left by three bits. Likewise, M := M / 2 * 3 shifts the bits to the right. Note that the compiler selects the appropriate shifting operator when translating these operations to machine code — no actual multiplication or division will be performed.

Let's see a simple implementation of the CRC-CCITT (0x1D0F) algorithm:

Listing 48: crc defs.ads

```
package Crc_Defs is
1
2
        type Byte is mod 2 ** 8;
3
        type Crc is mod 2 ** 16;
4
5
        type Byte_Array is
6
          array (Positive range <>) of Byte;
7
8
        function Crc_CCITT (A : Byte_Array)
9
                             return Crc;
10
11
        procedure Display (Crc_A : Crc);
12
13
        procedure Display (A : Byte_Array);
14
15
   end Crc_Defs;
16
```

1 2

3

7

Listing 49: crc_defs.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Crc_Defs is
3
4
        package Byte IO is new Modular IO (Byte);
5
        package Crc_I0 is new Modular_I0 (Crc);
6
7
        function Crc_CCITT (A : Byte_Array)
8
                              return Crc
9
10
        is
                  : Byte;
           Х
11
           Crc_A : Crc := 16#1d0f#;
12
        begin
13
           for I in A'Range loop
14
               X := Byte (Crc_A / 2 ** 8) xor A (I);
15
               X := X \text{ xor } (X / 2 ** 4);
16
               declare
17
                  Crc_X : constant Crc := Crc (X);
18
               begin
19
                  Crc_A := Crc_A * 2 ** 8 xor
20
                            Crc_X * 2 ** 12 xor
21
                            Crc_X * 2 ** 5 xor
22
                            Crc_X;
23
              end;
24
           end loop;
25
26
           return Crc A;
27
        end Crc_CCITT;
28
29
        procedure Display (Crc_A : Crc) is
30
        begin
31
           Crc_IO.Put (Crc_A);
32
           New_Line;
33
        end Display;
34
35
        procedure Display (A : Byte_Array) is
36
        begin
37
           for E of A loop
38
39
               Byte_IO.Put (E);
              Put (", ");
40
           end loop;
41
           New_Line;
42
        end Display;
43
44
   begin
45
       Byte_IO.Default_Width := 1;
46
       Byte IO.Default Base := 16;
47
       Crc_IO.Default_Width := 1;
48
       Crc_IO.Default_Base
                               := 16;
49
   end Crc_Defs;
50
```

Listing 50: show_crc.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Crc_Defs; use Crc_Defs;
procedure Show_Crc is
AA : constant Byte_Array :=
(16#0#, 16#20#, 16#30#);
```

```
Crc_A : Crc;
7
   begin
8
      Crc_A := Crc_CCITT (AA);
9
10
      Put ("Input array: ");
11
      Display (AA);
12
13
      Put ("CRC-CCITT: ");
14
      Display (Crc_A);
15
  end Show_Crc;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Modular_Types.Mod_Crc_CCITT_Ada
MD5: 9c66abfadcce92231295cbccad087912
```

Runtime output

```
Input array: 16#0#, 16#20#, 16#30#,
CRC-CCITT: 16#21B9#
```

In this example, the core of the algorithm is implemented in the Crc_CCITT function. There, we use bit shifting — for instance, $* \ 2 \ ** \ 8$ and $/ \ 2 \ ** \ 8$, which shift left and right, respectively, by eight bits. We also use the **xor** operator.

8.5 Attributes of Floating-Point Types

In this section, we discuss various attributes related to floating-point types.

```
1 In the Ada Reference Manual
```

- 3.5.8 Operations of Floating Point Types¹⁶⁰
- A.5.3 Attributes of Floating Point Types¹⁶¹

8.5.1 Representation-oriented attributes

In this section, we discuss attributes related to the representation of floating-point types.

Attribute: Machine_Radix

Machine_Radix is an attribute that returns the radix of the hardware representation of a type. For example:

Listing 51: show machine radix.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Show_Machine_Radix is
begin
Put_Line
("Float'Machine_Radix: "
  & Float'Machine_Radix'Image);
Put_Line
```

¹⁶⁰ http://www.ada-auth.org/standards/22rm/html/RM-3-5-8.html
¹⁶¹ http://www.ada-auth.org/standards/22rm/html/RM-A-5-3.html

```
9 ("Long_Float'Machine_Radix: "
10 & Long_Float'Machine_Radix'Image);
11 Put_Line
12 ("Long_Long_Float'Machine_Radix: "
13 & Long_Long_Float'Machine_Radix'Image);
14 end Show_Machine_Radix;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Machine_

⇔Radix

MD5: 88680df680f1db4ff803912850370551
```

Runtime output

Float'Machine_Radix: 2
Long_Float'Machine_Radix: 2
Long Long Float'Machine Radix: 2

Usually, this value is two, as the radix is based on a binary system.

Attributes: Machine_Mantissa

Machine_Mantissa is an attribute that returns the number of bits reserved for the mantissa of the floating-point type. For example:

Listing 52: show_machine_mantissa.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Machine Mantissa is
3
   begin
4
      Put Line
5
                                               n
        ("Float'Machine Mantissa:
6
         & Float'Machine_Mantissa'Image);
7
      Put Line
8
        ("Long_Float'Machine_Mantissa:
9
         & Long_Float'Machine_Mantissa'Image);
10
      Put_Line
11
        ("Long_Long_Float'Machine_Mantissa: "
12
         & Long Long Float'Machine Mantissa'Image);
13
  end Show_Machine_Mantissa;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Machine_

⊲Mantissa

MD5: da946a90a454c6e8f68cbff1ec54c7d3
```

Runtime output

Float'Machine_Mantissa:	24
Long_Float'Machine_Mantissa:	53
Long_Long_Float'Machine_Mantissa:	64

On a typical desktop PC, as indicated by Machine_Mantissa, we have 24 bits for the floatingpoint mantissa of the **Float** type.

Machine_Emin and Machine_Emax

The Machine_Emin and Machine_Emax attributes return the minimum and maximum value, respectively, of the machine exponent the floating-point type. Note that, in all cases, the returned value is a universal integer. For example:

```
Listing 53: show_machine_emin_emax.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show Machine Emin Emax is
3
   begin
4
      Put Line
5
         ("Float'Machine Emin:
6
         & Float 'Machine Emin'Image);
7
      Put Line
8
         ("Float'Machine_Emax:
9
         & Float'Machine_Emax'Image);
10
      Put Line
11
         ( "
           Long_Float'Machine_Emin:
12
         & Long_Float'Machine_Emin'Image);
13
      Put Line
14
         ("Long_Float'Machine Emax:
                                                n.
15
         & Long Float 'Machine Emax'Image);
16
      Put Line
17
         ("Long Long Float'Machine Emin:
18
          & Long Long Float 'Machine Emin'Image);
19
      Put Line
20
         ("Long Long Float'Machine Emax:
21
         & Long_Long_Float'Machine Emax'Image);
22
   end Show_Machine_Emin_Emax;
23
```

Code block metadata

Runtime output

<pre>Float'Machine_Emin:</pre>	-125
<pre>Float'Machine_Emax:</pre>	128
Long_Float'Machine_Emin:	-1021
Long_Float'Machine_Emax:	1024
Long_Long_Float'Machine_Emin:	-16381
Long_Long_Float'Machine_Emax:	16384

On a typical desktop PC, the value of **Float** 'Machine_Emin and **Float** 'Machine_Emax is -125 and 128, respectively.

To get the actual minimum and maximum value of the exponent for a specific type, we need to use the Machine_Radix attribute that we've seen previously. Let's calculate the minimum and maximum value of the exponent for the **Float** type on a typical PC:

- Value of minimum exponent: Float'Machine_Radix ** Float'Machine_Emin.
 - In our target platform, this is $2^{-125} = 2.35098870164457501594 \times 10^{-38}$.
- Value of maximum exponent: Float'Machine_Radix ** Float'Machine_Emax.
 - In our target platform, this is $2^{128} = 3.40282366920938463463 \times 10^{38}$.

Attribute: Digits

Digits is an attribute that returns the requested decimal precision of a floating-point subtype. Let's see an example:

```
Listing 54: show_digits.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Digits is
3
   begin
4
      Put Line ("Float'Digits:
5
                 & Float'Digits'Image);
6
      Put_Line ("Long_Float'Digits:
7
                 & Long_Float'Digits'Image);
8
      Put_Line ("Long_Long_Float'Digits: "
9
                 & Long_Long_Float'Digits'Image);
10
   end Show_Digits;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Digits
MD5: cd1c88054f7d54703760a852d08acb6d

Runtime output

Float'Digits: 6 Long_Float'Digits: 15 Long_Long_Float'Digits: 18

Here, the requested decimal precision of the Float type is six digits.

Note that we said that **Digits** is the *requested* level of precision, which is specified as part of declaring a floating point type. We can retrieve the actual decimal precision with Base'Digits. For example:

Listing 55: show base digits.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Base Digits is
3
      type Float_D3 is new Float digits 3;
4
   begin
5
                                               n
      Put Line ("Float D3'Digits:
6
                 & Float_D3'Digits'Image);
7
      Put Line ("Float D3'Base'Digits:
8
                 & Float_D3'Base'Digits'Image);
g
   end Show_Base_Digits;
10
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Base_Digits
MD5: a2deb352f93511ab2a39d41f0b3f9512

Runtime output

Float_D3'Digits: 3 Float_D3'Base'Digits: 6

The requested decimal precision of the Float_D3 type is three digits, while the actual decimal precision is six digits (on a typical desktop PC).

Attributes: Denorm, Signed_Zeros, Machine_Rounds, Machine_Overflows

In this section, we discuss attributes that return **Boolean** values indicating whether a feature is available or not in the target architecture:

- Denorm is an attribute that indicates whether the target architecture uses denormalized numbers¹⁶².
- Signed_Zeros is an attribute that indicates whether the type uses a sign for zero values, so it can represent both -0.0 and 0.0.
- Machine_Rounds is an attribute that indicates whether rounding-to-nearest is used, rather than some other choice (such as rounding-toward-zero).
- Machine_Overflows is an attribute that indicates whether a Constraint_Error exception is (or is not) guaranteed to be raised when an operation with that type produces an overflow or divide-by-zero.

Listing 56: show boolean attributes.ad
--

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Boolean Attributes is
3
   begin
4
      Put Line
5
         ("Float'Denorm:
6
          & Float 'Denorm 'Image);
7
      Put_Line
8
         ("Long_Float'Denorm:
9
          & Long Float 'Denorm 'Image);
10
      Put Line
11
         ("Long Long Float'Denorm: "
12
          & Long Long Float 'Denorm'Image);
13
       Put Line
14
         ("Float'Signed Zeros:
15
          & Float'Signed_Zeros'Image);
16
17
      Put Line
         ("Long_Float'Signed_Zeros:
18
          & Long_Float'Signed_Zeros'Image);
19
      Put Line
20
         ("Long Long Float'Signed Zeros: "
21
          & Long Long Float'Signed Zeros'Image);
22
       Put Line
23
         ("Float'Machine Rounds:
24
          & Float'Machine Rounds'Image);
25
      Put Line
26
         ("Long Float'Machine Rounds:
27
          & Long Float 'Machine Rounds' Image);
28
      Put Line
29
         ("Long_Long_Float'Machine_Rounds: "
30
          & Long_Long_Float 'Machine_Rounds'Image);
31
      Put Line
32
         ("Float'Machine Overflows:
33
          & Float 'Machine Overflows'Image);
34
       Put_Line
35
         ("Long Float'Machine Overflows:
36
          & Long_Float 'Machine_Overflows'Image);
37
       Put Line
38
           'Long_Long_Float'Machine_Overflows: "
         ( "
39
          & Long_Long_Float'Machine_Overflows'Image);
40
   end Show_Boolean_Attributes;
41
```

Code block metadata

¹⁶² https://en.wikipedia.org/wiki/Subnormal_number

Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Machine_ GRounds_Overflows MD5: b3f72c212cf00e697fe144a87eb72339

Runtime output

Float'Denorm:	TRUE		
Long_Float'Denorm:	TRUE		
Long_Long_Float'Denorm:	TRUE		
Float'Signed_Zeros:		TRUE	
Long_Float'Signed_Zeros	:	TRUE	
Long_Long_Float'Signed_2	Zeros:	TRUE	
Float'Machine_Rounds:		TRI	JE
Long_Float'Machine_Round	ds:	TRI	JE
Long_Long_Float'Machine_Rounds: TRUE			
Float'Machine_Overflows	:		FALSE
Long_Float'Machine_Over	flows:		FALSE
Long_Long_Float'Machine	_Overf	lows:	FALSE

On a typical PC, we have the following information:

- Denorm is true (i.e. the architecture uses denormalized numbers);
- Signed_Zeros is true (i.e. the standard floating-point types use a sign for zero values);
- Machine_Rounds is true (i.e. rounding-to-nearest is used for floating-point types);
- Machine_Overflows is false (i.e. there's no guarantee that a Constraint_Error exception is raised when an operation with a floating-point type produces an overflow or divide-by-zero).

8.5.2 Primitive function attributes

In this section, we discuss attributes that we can use to manipulate floating-point values.

Attributes: Fraction, Exponent and Compose

The Exponent and Fraction attributes return "parts" of a floating-point value:

- Exponent returns the machine exponent, and
- Fraction returns the mantissa part.

Compose is used to return a floating-point value based on a fraction (the mantissa part) and the machine exponent.

Let's see some examples:

```
Listing 57: show_exponent_fraction_compose.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Exponent Fraction Compose is
3
   begin
4
      Put_Line
5
        ("Float'Fraction (1.0):
6
         & Float'Fraction (1.0)'Image);
7
      Put_Line
8
        ("Float'Fraction (0.25):
                                      п
9
         & Float'Fraction (0.25)'Image);
10
      Put Line
11
        ("Float'Fraction (1.0e-25): "
12
         & Float'Fraction (1.0e-25)'Image);
13
```

```
Put Line
14
         ("Float'Exponent (1.0):
15
         & Float'Exponent (1.0)'Image);
16
      Put Line
17
         ("Float'Exponent (0.25):
18
         & Float'Exponent (0.25)'Image);
19
      Put_Line
20
         ("Float'Exponent (1.0e-25): "
21
         & Float'Exponent (1.0e-25)'Image);
22
      Put Line
23
         ("Float'Compose (5.00000e-01, 1):
24
         & Float 'Compose (5.00000e-01, 1) 'Image);
25
26
      Put Line
         ("Float'Compose (5.00000e-01, -1): "
27
         & Float 'Compose (5.00000e-01, -1) 'Image);
28
      Put Line
29
         ("Float'Compose (9.67141E-01, -83): "
30
         & Float 'Compose (9.67141E-01, -83) 'Image);
31
   end Show_Exponent_Fraction_Compose;
32
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Exponent_

⊸Fraction

MD5: d2e61b6b9a7a50861145f6b65e9fac39
```

Runtime output

```
Float'Fraction (1.0): 5.00000E-01
Float'Fraction (0.25): 5.00000E-01
Float'Fraction (1.0e-25): 9.67141E-01
Float'Exponent (1.0): 1
Float'Exponent (0.25): -1
Float'Exponent (1.0e-25): -83
Float'Compose (5.00000e-01, 1): 1.00000E+00
Float'Compose (5.00000e-01, -1): 2.50000E-01
Float'Compose (9.67141E-01, -83): 1.00000E-25
```

To understand this code example, we have to take this formula into account:

Value = Fraction x Machine_Radix^{Exponent}

Considering that the value of **Float** 'Machine_Radix on a typical PC is two, we see that the value 1.0 is composed by a fraction of 0.5 and a machine exponent of one. In other words:

 $0.5 \times 2^1 = 1.0$

For the value 0.25, we get a fraction of 0.5 and a machine exponent of -1, which is the result of $0.5 \times 2^{-1} = 0.25$. We can use the Compose attribute to perform this calculation. For example, **Float** 'Compose (0.5, -1) = 0.25.

Note that Fraction is always between 0.5 and 0.999999 (i.e < 1.0), except for denormalized numbers, where it can be < 0.5.

Attribute: Scaling

Scaling is an attribute that scales a floating-point value based on the machine radix and a machine exponent passed to the function. For example:

Listing 58: show_scaling.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Scaling is
3
   begin
4
      Put Line ("Float'Scaling (0.25, 1): "
5
                 & Float'Scaling (0.25, 1)'Image);
6
      Put Line ("Float'Scaling (0.25, 2): "
7
                 & Float'Scaling (0.25, 2)'Image);
8
      Put_Line ("Float'Scaling (0.25, 3): "
9
                 & Float'Scaling (0.25, 3)'Image);
10
   end Show_Scaling;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Scaling
MD5: 9fa821d32911b74ee4b4fde3f3adafd8

Runtime output

Float'Scaling (0.25, 1): 5.00000E-01 Float'Scaling (0.25, 2): 1.00000E+00 Float'Scaling (0.25, 3): 2.00000E+00

The scaling is calculated with this formula:

scaling = value x Machine Radix^{machine exponent}

For example, on a typical PC with a machine radix of two, **Float** 'Scaling (0.25, 3) = 2.0 corresponds to

 $0.25 \times 2^3 = 2.0$

Round-up and round-down attributes

Floor and Ceiling are attributes that returned the rounded-down or rounded-up value, respectively, of a floating-point value. For example:

Listing 59: show_floor_ceiling.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Floor Ceiling is
3
   begin
4
      Put Line ("Float'Floor (0.25):
5
                & Float'Floor (0.25)'Image);
6
      Put Line ("Float'Ceiling (0.25): "
7
                & Float'Ceiling (0.25)'Image);
Q
  end Show_Floor_Ceiling;
a
```

Code block metadata

Runtime output

```
Float'Floor (0.25): 0.0000E+00
Float'Ceiling (0.25): 1.0000E+00
```

As we can see in this example, the rounded-down value (floor) of 0.25 is 0.0, while the rounded-up value (ceiling) of 0.25 is 1.0.

Round-to-nearest attributes

In this section, we discuss three attributes used for rounding: Rounding, Unbiased_Rounding, Machine_Rounding In all cases, the rounding attributes return the nearest integer value (as a floating-point value). For example, the rounded value for 4.8 is 5.0 because 5 is the closest integer value.

Let's see a code example:

Listing 60: show_	roundings.adb
-------------------	---------------

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Roundings is
3
4
   begin
      Put_Line
5
         ("Float'Rounding (0.5):
                                   .0
6
         & Float'Rounding (0.5)'Image);
7
      Put Line
8
         ("Float'Rounding (1.5): "
9
         & Float'Rounding (1.5)'Image);
10
      Put Line
11
         ("Float'Rounding (4.5): "
12
         & Float'Rounding (4.5)'Image);
13
      Put Line
14
         ("Float'Rounding (-4.5): "
15
         & Float'Rounding (-4.5)'Image);
16
      Put Line
17
         ("Float'Unbiased Rounding (0.5): "
18
         & Float'Unbiased_Rounding (0.5)'Image);
19
      Put Line
20
         ("Float'Unbiased Rounding (1.5): "
21
          & Float'Unbiased_Rounding (1.5)'Image);
22
      Put Line
23
         ("Float'Machine_Rounding (0.5): "
24
         & Float 'Machine Rounding (0.5) 'Image);
25
      Put Line
26
         ("Float'Machine_Rounding (1.5): "
27
         & Float 'Machine Rounding (1.5) 'Image);
28
   end Show Roundings;
29
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Rounding
MD5: 3f78165f092a163339cb9593ff15a50d

Runtime output

Float'Rounding (0.5): 1.00000E+00 Float'Rounding (1.5): 2.00000E+00 Float'Rounding (4.5): 5.00000E+00 Float'Rounding (-4.5): -5.00000E+00 Float'Unbiased_Rounding (0.5): 0.00000E+00 Float'Unbiased_Rounding (1.5): 2.00000E+00 Float'Machine_Rounding (1.5): 2.00000E+00

The difference between these attributes is the way they handle the case when a value is exactly in between two integer values. For example, 4.5 could be rounded up to 5.0 or rounded down to 4.0. This is the way each rounding attribute works in this case:

- Rounding rounds away from zero. Positive floating-point values are rounded up, while negative floating-point values are rounded down when the value is between two integer values. For example:
 - 4.5 is rounded-up to 5.0, i.e. Float'Rounding (4.5) = Float'Ceiling (4.5) = 5.0.
 - -4.5 is rounded-down to -5.0, i.e. Float'Rounding (-4.5) = Float'Floor (-4. 5) = -5.0.
- Unbiased_Rounding rounds toward the even integer. For example,
 - Float 'Unbiased_Rounding (0.5) = 0.0 because zero is the closest even integer, while
 - Float 'Unbiased_Rounding (1.5) = 2.0 because two is the closest even integer.
- Machine_Rounding uses the most appropriate rounding instruction available on the target platform. While this rounding attribute can potentially have the best performance, its result may be non-portable. For example, whether the rounding of 4.5 becomes 4.0 or 5.0 depends on the target platform.
 - If an algorithm depends on a specific rounding behavior, it's best to avoid the Machine_Rounding attribute. On the other hand, if the rounding behavior won't have a significant impact on the results, we can safely use this attribute.

Attributes: Truncation, Remainder, Adjacent

The Truncation attribute returns the *truncated* value of a floating-point value, i.e. the value corresponding to the integer part of a number rounded toward zero. This corresponds to the number before the radix point. For example, the truncation of 1.55 is 1.0 because the integer part of 1.55 is 1.

The Remainder attribute returns the remainder part of a division. For example, **Float** 'Remainder (1.25, 0.5) = 0.25. Let's briefly discuss the details of this operations. The result of the division 1.25 / 0.5 is 2.5. Here, 1.25 is the dividend and 0.5 is the divisor. The quotient and remainder of this division are 2 and 0.25, respectively. (Here, the quotient is an integer number, and the remainder is the floating-point part that remains.)

Note that the relation between quotient and remainder is defined in such a way that we get the original dividend back when we use the formula: "quotient x divisor + remainder = dividend". For the previous example, this means $2 \times 0.5 + 0.25 = 1.25$.

The Adjacent attribute is the next machine value towards another value. For example, on a typical PC, the adjacent value of a small value — say, 1.0×10^{-83} — towards zero is +0.0, while the adjacent value of this small value towards 1.0 is another small, but greater value — in fact, it's 1.40130 x 10^{-45} . Note that the first parameter of the Adjacent attribute is the value we want to analyze and the second parameter is the Towards value.

Let's see a code example:

```
Listing 61: show_truncation_remainder_adjacent.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Truncation Remainder Adjacent is
3
   begin
4
      Put Line
5
        ("Float'Truncation (1.55): "
6
         & Float'Truncation (1.55)'Image);
7
      Put Line
8
        ("Float'Truncation (-1.55): "
9
         & Float'Truncation (-1.55)'Image);
10
      Put Line
11
```

```
("Float'Remainder (1.25, 0.25): "
12
         & Float'Remainder (1.25, 0.25)'Image);
13
      Put Line
14
         ("Float'Remainder (1.25, 0.5):
                                          н
15
         & Float'Remainder (1.25, 0.5)'Image);
16
      Put_Line
17
         ("Float'Remainder (1.25, 1.0):
18
         & Float'Remainder (1.25, 1.0)'Image);
19
      Put Line
20
         ("Float'Remainder (1.25, 2.0): "
21
         & Float'Remainder (1.25, 2.0)'Image);
22
      Put Line
23
         ("Float'Adjacent (1.0e-83, 0.0): "
24
         & Float'Adjacent (1.0e-83, 0.0)'Image);
25
      Put Line
26
         ("Float'Adjacent (1.0e-83, 1.0): "
27
         & Float'Adjacent (1.0e-83, 1.0)'Image);
28
   end Show_Truncation_Remainder_Adjacent;
29
```

Attributes: Copy_Sign and Leading_Part

Copy_Sign is an attribute that returns a value where the sign of the second floating-point argument is multiplied by the magnitude of the first floating-point argument. For example, **Float**'Copy_Sign (1.0, -10.0) is -1.0. Here, the sign of the second argument (-10.0) is multiplied by the magnitude of the first argument (1.0), so the result is -1.0.

Leading_Part is an attribute that returns the *approximated* version of the mantissa of a value based on the specified number of leading bits for the mantissa. Let's see some examples:

- Float 'Leading_Part (3.1416, 1) is 2.0 because that's the value we can represent with one leading bit.
 - Note that **Float**'Fraction (2.0) = 0.5 which can be represented with one leading bit in the mantissa and **Float**'Exponent (2.0) = 2.)
- If we increase the number of leading bits of the mantissa to two by writing **Float**'Leading_Part (3.1416, 2) —, we get 3.0 because that's the value we can represent with two leading bits.
- If we increase again the number of leading bits to five Float'Leading_Part (3. 1416, 5) —, we get 3.125.
 - Note that, in this case Float 'Fraction (3.125) = 0.78125 and Float 'Exponent (3.125) = 2.
 - The binary mantissa is actually 2#110_0100_0000_0000_0000_0000#, which can be represented with five leading bits as expected: 2#110_01#.
 - * We can get the binary mantissa by calculating Float 'Fraction (3.125) * Float (Float 'Machine_Radix) ** (Float 'Machine_Mantissa - 1) and converting the result to binary format. The -1 value in the formula corresponds to the sign bit.

1 Attention

In this explanation about the Leading_Part attribute, we're talking about leading bits. Strictly speaking, however, this is actually a simplification, and it's only correct if Machine_Radix is equal to two — which is the case for most machines. Therefore, in most cases, the explanation above is perfectly acceptable.

However, if Machine_Radix is *not* equal to two, we cannot use the term "bits" anymore, but rather digits of the Machine_Radix.

Let's see some examples:

Listing 62: show_copy_sign_leading_part_machine.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Copy Sign Leading Part Machine is
3
   begin
4
       Put_Line
5
         ("Float'Copy Sign (1.0, -10.0): "
6
          & Float'Copy Sign (1.0, -10.0)'Image);
7
       Put Line
8
         ("Float'Copy_Sign (-1.0, -10.0): "
9
          & Float'Copy Sign (-1.0, -10.0)'Image);
10
       Put Line
11
         ("Float'Copy_Sign (1.0, 10.0): "
   & Float'Copy_Sign (1.0, 10.0)'Image);
12
13
       Put Line
14
         ("Float'Copy_Sign (1.0, -0.0):
                                            н
15
          & Float'Copy_Sign (1.0, -0.0)'Image);
16
       Put Line
17
         ("Float'Copy_Sign (1.0, 0.0): "
18
          & Float'Copy_Sign (1.0, 0.0)'Image);
19
       Put Line
20
         ("Float'Leading Part (1.75, 1): "
21
          & Float'Leading Part (1.75, 1)'Image);
22
       Put Line
23
         ("Float'Leading Part (1.75, 2): "
24
          & Float'Leading Part (1.75, 2)'Image);
25
       Put Line
26
         ("Float'Leading Part (1.75, 3): "
27
          & Float'Leading Part (1.75, 3)'Image);
28
   end Show Copy_Sign_Leading_Part_Machine;
29
```

Attribute: Machine

Not every real number is directly representable as a floating-point value on a specific machine. For example, let's take a value such as 1.0×10^{15} (or 1,000,000,000,000):

Listing 63: show_float_value.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Float Value is
3
      package F IO is new
4
         Ada.Text IO.Float IO (Float);
5
6
      V : Float;
7
   begin
8
      F IO.Default Fore := 3;
9
      F_IO.Default_Aft := 1;
10
      F_IO.Default_Exp := 0;
11
12
      V := 1.0E+15;
13
      Put ("1.0E+15 = ");
14
      F IO.Put (Item => V);
15
      New Line;
16
```

17
18 end Show_Float_Value;

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Floating_Point_Types.Float_Value
MD5: a7f80f7584ebaf39f2d5f9564c9c7d64

Runtime output

1.0E+15 = 999999986991000.0

If we run this example on a typical PC, we see that the expected value 1_{000}_{00

This *automatic* modification we've just described is actually hidden, so to say, in the assignment. However, we can make it more visible by using the Machine (X) attribute, which returns a version of X that is representable on the target machine. The Machine (X) attribute rounds (or truncates) X to either one of the adjacent machine numbers for the specific floating-point type of X. (Of course, if the real value of X is directly representable on the target machine, no modification is performed.)

In fact, we could rewrite the V := 1.0E+15 assignment of the code example as V := **Float** 'Machine (1.0E+15), as we're never assigning a real value directly to a floatingpointing variable — instead, we're first converting it to a version of the real value that is representable on the target machine. In this case, 999999986991000.0 is a representable version of the real value 1.0×10^{15} . Of course, writing V := 1.0E+15 or V := **Float** 'Machine (1.0E+15) doesn't make any difference to the actual value that is assigned to V (in the case of this specific target architecture), as the conversion to a representable value happens automatically during the assignment to V.

There are, however, instances where using the Machine attribute does make a difference in the result. For example, let's say we want to calculate the difference between the original real value in our example (1.0×10^{15}) and the actual value that is assigned to V. We can do this by using the Machine attribute in the calculation:

Listing 6	54:	show	machine	attribute.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Machine Attribute is
3
      package F_IO is new
4
        Ada.Text_I0.Float_I0 (Float);
5
6
      V : Float;
7
   begin
8
      F IO.Default Fore := 3;
9
      F IO.Default Aft := 1;
10
      F IO.Default Exp := 0;
11
12
13
      Put Line
         ("Original value: 1_000_000_000_000_000.0");
14
15
      V := 1.0E + 15;
16
      Put ("Machine value: ");
17
      F_IO.Put (Item => V);
18
      New_Line;
19
20
```

```
21 V := 1.0E+15 - Float'Machine (1.0E+15);
22 Put ("Difference: ");
23 F_I0.Put (Item => V);
24 New_Line;
25
26 end Show_Machine_Attribute;
```

Code block metadata

Runtime output

Original value: 1_000_000_000_000_000.0 Machine value: 99999986991000.0 Difference: 13008896.0

When we run this example on a typical PC, we see that the difference is roughly 1.3009 x 10^7 . (Actually, the value that we might see is 1.3008896 x 10^7 , which is a version of 1.3009 x 10^7 that is representable on the target machine.)

When we write 1.0E+15 - Float Machine (1.0E+15):

- the first value in the operation is the universal real value 1.0×10^{15} , while
- the second value in the operation is a version of the universal real value 1.0×10^{15} that is representable on the target machine.

This also means that, in the assignment to V, we're actually writing V := Float 'Machine (1.0E+15 - Float 'Machine (1.0E+15)), so that:

- 1. we first get the intermediate real value that represents the difference between these values; and then
- 2. we get a version of this intermediate real value that is representable on the target machine.

This is the reason why we see 1.3008896 x 10^7 instead of 1.3009 x 10^7 when we run this application.

8.6 Attributes of Fixed-Point types

In this section, we discuss various attributes and operations related to fixed-point types.

```
1 In the Ada Reference Manual
```

- 3.5.10 Operations of Fixed Point Types¹⁶³
- A.5.4 Attributes of Fixed Point Types¹⁶⁴

8.6.1 Attributes of ordinary and decimal fixed-point types

¹⁶³ http://www.ada-auth.org/standards/22rm/html/RM-3-5-10.html

¹⁶⁴ http://www.ada-auth.org/standards/22rm/html/RM-A-5-4.html

Attribute: Machine_Radix

Machine_Radix is an attribute that returns the radix of the hardware representation of a type. For example:

```
Listing 65: show_fixed_machine_radix.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Fixed Machine Radix is
3
      type T3_D3 is delta 10.0 ** (-3) digits 3;
4
5
      D : constant := 2.0 ** (-31);
6
      type TQ31 is delta D range -1.0 .. 1.0 - D;
7
   begin
8
      Put_Line ("T3_D3'Machine_Radix: "
9
                 & T3_D3'Machine_Radix'Image);
10
      Put_Line ("TQ31'Machine_Radix: "
11
                 & TQ31'Machine_Radix'Image);
12
   end Show Fixed Machine Radix;
13
```

Code block metadata

Runtime output

T3_D3'Machine_Radix: 2 TQ31'Machine_Radix: 2

Usually, this value is two, as the radix is based on a binary system.

Attribute: Machine_Rounds and Machine_Overflows

In this section, we discuss attributes that return **Boolean** values indicating whether a feature is available or not in the target architecture:

- Machine_Rounds is an attribute that indicates what happens when the result of a fixedpoint operation is inexact:
 - T'Machine_Rounds = True: inexact result is rounded;
 - T'Machine_Rounds = False: inexact result is truncated.
- Machine_Overflows is an attribute that indicates whether a Constraint_Error is guaranteed to be raised when a fixed-point operation with that type produces an overflow or divide-by-zero.

Listing 66: show_boolean_attributes.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Boolean Attributes is
3
      type T3_D3 is delta 10.0 ** (-3) digits 3;
4
5
      D : constant := 2.0 ** (-31);
6
      type TQ31 is delta D range -1.0 .. 1.0 - D;
7
8
   begin
                                            п
      Put Line ("T3 D3'Machine Rounds:
9
                 & T3_D3'Machine_Rounds'Image);
10
      Put Line ("TQ31'Machine Rounds:
11
```

```
12 & TQ31'Machine_Rounds'Image);
13 Put_Line ("T3_D3'Machine_Overflows: "
14 & T3_D3'Machine_Overflows'Image);
15 Put_Line ("TQ31'Machine_Overflows: "
16 & TQ31'Machine_Overflows'Image);
17 end Show_Boolean_Attributes;
```

Attribute: Small and Delta

The Small and **Delta** attributes return numbers that indicate the numeric precision of a fixed-point type. In many cases, the Small of a type T is equal to the **Delta** of that type — i.e. T'Small = T'Delta. Let's discuss each attribute and how they distinguish from each other.

The **Delta** attribute returns the value of the **delta** that was used in the type definition. For example, if we declare **type T3_D3 is delta** 10.0 ** (-3) **digits** D, then the value of T3_D3'Delta is the 10.0 ** (-3) that we used in the type definition.

The Small attribute returns the "small" of a type, i.e. the smallest value used in the machine representation of the type. The *small* must be at least equal to or smaller than the *delta* — in other words, it must conform to the T'Small <= T'Delta rule.

For further reading...

The Small and the **Delta** need not actually be small numbers. They can be arbitrarily large. For instance, they could be 1.0, or 1000.0. Consider the following example:

```
Listing 67: fixed point defs.ads
```

```
package Fixed Point Defs is
1
      S
            : constant := 32;
2
      Exp
            : constant := 128;
3
            : constant := 2.0 ** (-S + Exp + 1);
4
      D
5
      type Fixed is delta D
6
        range -1.0 * 2.0 ** Exp ..
7
               1.0 * 2.0 ** Exp - D;
8
9
      pragma Assert (Fixed'Size = S);
10
   end Fixed Point Defs;
11
```

```
Listing 68: show fixed type info.adb
```

```
with Fixed Point Defs; use Fixed Point Defs;
1
   with Ada.Text IO;
                            use Ada.Text I0;
2
3
   procedure Show_Fixed_Type_Info is
4
   begin
5
      Put_Line ("Size : "
6
                 & Fixed'Size'Image);
7
      Put Line ("Small : "
8
                 & Fixed'Small'Image);
      Put Line ("Delta : "
10
                 & Fixed'Delta'Image);
11
      Put Line ("First : "
12
                 & Fixed'First'Image);
13
      Put Line ("Last :
14
                 & Fixed'Last'Image);
15
   end Show_Fixed_Type_Info;
16
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Fixed_Point_Types.Large_Small_ Attribute

MD5: 89672950b355060d250e0f5d7e2d40cb

Runtime output

Size : 32

9

Small : 1.58456325028528675E+29 Delta : 1.58456325028528675E+29 First : -340282366920938463463374607431768211456.0 Last : 340282366762482138434845932244680310784.0

In this example, the small of the Fixed type is actually quite large: 1.58456325028528675^{29} . (Also, the first and the last values are -340,282,366,920,938,463,463,374,607,431,768,211,456.0 large: and 340,282,366,762,482,138,434,845,932,244,680,310,784.0, or approximately -3.4028³⁸ and 3.4028³⁸.)

In this case, if we assign 1 or 1,000 to a variable F of this type, the actual value stored in F is zero. Feel free to try this out!

When we declare an ordinary fixed-point data type, we must specify the *delta*. Specifying the *small*, however, is optional:

- If the small isn't specified, it is automatically selected by the compiler. In this case, the actual value of the *small* is an implementation-defined power of two — always following the rule that says: T'Small <= T'Delta.
- If we want, however, to specify the *small*, we can do that by using the Small aspect. In this case, it doesn't need to be a power of two.

For decimal fixed-point types, we cannot specify the *small*. In this case, it's automatically selected by the compiler, and it's always equal to the *delta*.

Let's see an example:

Listing 69: fixed small delta.ads

```
package Fixed_Small_Delta is
1
      D3 : constant := 10.0 ** (-3);
2
3
      type T3 D3 is delta D3 digits 3;
4
```

```
type TD3
                  is delta D3 range -1.0 .. 1.0 - D3;
6
7
      D31 : constant := 2.0 ** (-31);
8
      D15 : constant := 2.0 ** (-15);
9
10
      type TQ31 is delta D31 range -1.0 .. 1.0 - D31;
11
12
      type TQ15 is delta D15 range -1.0 .. 1.0 - D15
13
        with Small => D31;
14
   end Fixed_Small_Delta;
15
```

5

Listing 70: show_fixed_small_delta.adb

```
use Ada.Text I0;
   with Ada.Text IO;
1
2
   with Fixed Small Delta; use Fixed Small Delta;
3
4
   procedure Show_Fixed_Small_Delta is
5
   begin
6
      Put_Line ("T3_D3'Small: "
7
                & T3_D3'Small'Image);
8
      Put_Line ("T3_D3'Delta:
9
                & T3 D3'Delta'Image);
10
      Put Line ("T3 D3'Size: "
11
                & T3_D3'Size'Image);
12
      Put Line ("-----");
13
14
      Put_Line ("TD3'Small: "
15
                & TD3'Small'Image);
16
      Put_Line ("TD3'Delta: "
17
                & TD3'Delta'Image);
18
      Put_Line ("TD3'Size: "
19
                & TD3'Size'Image);
20
      Put_Line ("-----");
21
22
      Put_Line ("TQ31'Small: "
23
                & TQ31'Small'Image);
24
      Put Line ("TQ31'Delta: "
25
                & TQ31'Delta'Image);
26
      Put_Line ("TQ32'Size: "
27
                & TQ31'Size'Image);
28
      Put Line ("-----");
29
30
      Put_Line ("TQ15'Small: "
31
                & TQ15'Small'Image);
32
      Put_Line ("TQ15'Delta: "
33
                & TQ15'Delta'Image);
34
      Put_Line ("TQ15'Size: "
35
                & TQ15'Size'Image);
36
   end Show_Fixed_Small_Delta;
37
```

Code block metadata

Runtime output

As we can see in the output of the code example, the **Delta** attribute returns the value we used for **delta** in the type definition of the T3_D3, TD3, TQ31 and TQ15 types.

The TD3 type is an ordinary fixed-point type with the the same delta as the decimal T3_D3 type. In this case, however, TD3'Small is not the same as the TD3'Delta. On a typical desktop PC, TD3'Small is 2⁻¹⁰, while the delta is 10⁻³. (Remember that, for ordinary fixed-point types, if we don't specify the *small*, it's automatically selected by the compiler as a power of two smaller than or equal to the *delta*.)

In the case of the TQ15 type, we're specifying the *small* by using the Small aspect. In this case, the underlying size of the TQ15 type is 32 bits, while the precision we get when operating with this type is 16 bits. Let's see a specific example for this type:

Listing 71: show_fixed_small_delta.adb

```
with Ada.Text IO;
                             use Ada.Text I0;
1
2
   with Fixed Small Delta; use Fixed Small Delta;
3
4
   procedure Show Fixed Small Delta is
5
      V : TQ15;
6
   beain
7
      Put Line ("V'Size: " & V'Size'Image);
8
9
      V := TQ15'Small;
10
      Put Line ("V: " & V'Image);
11
12
      V := TQ15'Delta;
13
      Put Line ("V: " & V'Image);
14
   end Show Fixed Small Delta;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Fixed_Point_Types.Fixed_Small_

→Delta

MD5: f2a71db911913d6fbf5343671599c0ae
```

Runtime output

V'Size: 32 V: 0.00000 V: 0.00003

In the first assignment, we assign TQ15'Small (2^{-31}) to V. This value is smaller than the type's *delta* (2^{-15}) . Even though V'Size is 32 bits, V'Delta indicates 16-bit precision, and TQ15'Small requires 32-bit precision to be represented correctly. As a result, V has a value

of zero after this assignment.

In contrast, after the second assignment — where we assign TQ15'Delta (2^{-15}) to V — we see, as expected, that V has the same value as the *delta*.

Attributes: Fore and Aft

The Fore and Aft attributes indicate the number of characters or digits needed for displaying a value in decimal representation. To be more precise:

- The Fore attribute refers to the digits before the decimal point and it returns the number of digits plus one for the sign indicator (which is either - or space), and it's always at least two.
- The Aft attribute returns the number of decimal digits that is needed to represent the delta after the decimal point.

Let's see an example:

```
Listing 72: show_fixed_fore_aft.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Fixed Fore Aft is
3
      type T3_D3 is delta 10.0 ** (-3) digits 3;
4
5
      D : constant := 2.0 ** (-31);
6
      type TQ31 is delta D range -1.0 .. 1.0 - D;
7
8
      Dec : constant T3 D3 := -0.123;
9
      Fix : constant TQ31 := -TQ31'Delta;
10
   begin
11
      Put_Line ("T3_D3'Fore: "
12
                 & T3 D3'Fore'Image);
13
      Put_Line ("T3 D3'Aft:
14
                 & T3 D3'Aft'Image);
15
16
      Put Line ("TQ31'Fore: "
17
                 & TQ31'Fore'Image);
18
      Put_Line ("TQ31'Aft: "
19
                 & TQ31'Aft'Image);
20
      Put Line ("----");
21
      Put Line ("Dec: "
22
                 & Dec'Image);
23
      Put Line ("Fix: "
24
                 & Fix'Image);
25
   end Show Fixed Fore Aft;
26
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Fixed_Point_Types.Fixed_Fore_Aft
MD5: d031f74d967a96dee1c6a83ff4bd14cf

Runtime output

T3_D3'Fore: 2 T3_D3'Aft: 3 TQ31'Fore: 2 TQ31'Aft: 10 ----Dec: -0.123 Fix: -0.0000000005 As we can see in the output of the Dec and Fix variables at the bottom, the value of Fore is two for both T3_D3 and TQ31. This value corresponds to the length of the string "-0" displayed in the output for these variables (the first two characters of "-0.123" and "-0.0000000005").

The value of Dec'Aft is three, which matches the number of digits after the decimal point in "-0.123". Similarly, the value of Fix'Aft is 10, which matches the number of digits after the decimal point in "-0.0000000005".

8.6.2 Attributes of decimal fixed-point types

The attributes presented in this subsection are only available for decimal fixed-point types.

Attribute: Digits

Digits is an attribute that returns the number of significant decimal digits of a decimal fixed-point subtype. This corresponds to the value that we use for the **digits** in the definition of a decimal fixed-point type.

Let's see an example:

Listing 73: show_decimal_digits.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Decimal_Digits is
3
      type T3 D6 is delta 10.0 ** (-3) digits 6;
4
      subtype T3_D2 is T3_D6 digits 2;
5
   begin
6
      Put_Line ("T3_D6'Digits: "
7
                 & T3 D6'Digits'Image);
8
      Put Line ("T3 D2'Digits: "
9
                 & T3_D2'Digits'Image);
10
   end Show_Decimal_Digits;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Fixed_Point_Types.Decimal_Digits
MD5: d46e67bd0f8b369918e7ab9ab4413ae7

Runtime output

T3_D6'Digits: 6 T3_D2'Digits: 2

In this example, T3_D6'Digits is six, which matches the value that we used for **digits** in the type definition of T3_D6. The same logic applies for subtypes, as we can see in the value of T3_D2'Digits. Here, the value is two, which was used in the declaration of the T3_D2 subtype.

Attribute: Scale

According to the Ada Reference Manual, the Scale attribute "indicates the position of the point relative to the rightmost significant digits of values" of a decimal type. For example:

- If the value of Scale is two, then there are two decimal digits after the decimal point.
- If the value of Scale is negative, that implies that the **Delta** is a power of 10 greater than 1, and it would be the number of zero digits that every value would end in.

The Scale corresponds to the N used in the **delta** 10.0 ** (-N) expression of the type declaration. For example, if we write **delta** 10.0 ** (-3) in the declaration of a type T, then the value of T'Scale is three.

Let's look at this complete example:

```
Listing 74: show_decimal_scale.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
   procedure Show Decimal Scale is
3
      type TM3_D6 is delta 10.0 **
                                       3 digits 6;
4
      type T3_D6 is delta 10.0 ** (-3) digits 6;
5
      type T9_D12 is delta 10.0 ** (-9) digits 12;
6
   begin
7
      Put Line ("TM3 D6'Scale: "
8
                 & TM3 D6'Scale'Image);
9
      Put Line ("T3 D6'Scale: "
10
                 & T3 D6'Scale'Image);
11
      Put Line ("T9 D12'Scale: "
12
                 & T9 D12'Scale'Image);
13
   end Show_Decimal_Scale;
14
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Fixed_Point_Types.Decimal_Scale
MD5: 56a99848cf31a9c69fe6d91ead73375a

Runtime output

```
TM3_D6'Scale: -3
T3_D6'Scale: 3
T9_D12'Scale: 9
```

In this example, we get the following values for the scales:

- TM3_D6'Scale = -3,
- T3_D6'Scale = 3,
- T9 D12 = 9.

As you can see, the value of Scale is directly related to the *delta* of the corresponding type declaration.

Attribute: Round

The Round attribute rounds a value of any real type to the nearest value that is a multiple of the *delta* of the decimal fixed-point type, rounding away from zero if exactly between two such multiples.

For example, if we have a type T with three digits, and we use a value with 10 digits after the decimal point in a call to T'Round, the resulting value will have three digits after the decimal point.

Note that the X input of an S'Round (X) call is a universal real value, while the returned value is of S'Base type.

Let's look at this example:

Listing 75: show_decimal_round.adb

```
with Ada.Text_I0; use Ada.Text_I0;
```

```
procedure Show_Decimal_Round is
3
      type T3_D3 is delta 10.0 ** (-3) digits 3;
4
   begin
5
      Put_Line ("T3_D3'Round (0.2774): "
6
                 & T3_D3'Round (0.2774)'Image);
7
      Put_Line ("T3_D3'Round (0.2777): "
8
                 & T3_D3'Round (0.2777)'Image);
9
   end Show_Decimal_Round;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Fixed_Point_Types.Decimal_Round
MD5: 153d9dae52fee750da30dd9152a03c37
```

Runtime output

```
T3_D3'Round (0.2774): 0.277
T3_D3'Round (0.2777): 0.278
```

Here, the T3_D3 has a precision of three digits. Therefore, to fit this precision, 0.2774 is rounded to 0.277, and 0.2777 is rounded to 0.278.

8.7 Big Numbers

As we've seen before, we can define numeric types in Ada with a high degree of precision. However, these normal numeric types in Ada are limited to what the underlying hardware actually supports. For example, any signed integer type — whether defined by the language or the user — cannot have a range greater than that of System.Min_Int ... System. Max_Int because those constants reflect the actual hardware's signed integer types. In certain applications, that precision might not be enough, so we have to rely on arbitraryprecision arithmetic¹⁶⁵. These so-called "big numbers" are limited conceptually only by available memory, in contrast to the underlying hardware-defined numeric types.

Ada supports two categories of big numbers: big integers and big reals — both are specified in child packages of the Ada.Numerics.Big_Numbers package:

Category	Package
Big Integers	Ada.Numerics.Big_Numbers.Big_Integers
Big Reals	Ada.Numerics.Big_Numbers.Big_Real

In the Ada Reference Manual

- Big Numbers¹⁶⁶
- Big Integers¹⁶⁷
- Big Reals¹⁶⁸

¹⁶⁵ https://en.wikipedia.org/wiki/arbitrary-precision_arithmetic

¹⁶⁶ http://www.ada-auth.org/standards/22rm/html/RM-A-5-5.html

¹⁶⁷ http://www.ada-auth.org/standards/22rm/html/RM-A-5-6.html

¹⁶⁸ http://www.ada-auth.org/standards/22rm/html/RM-A-5-7.html

8.7.1 Overview

Let's start with a simple declaration of big numbers:

```
Listing 76: show_simple_big_numbers.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Numerics.Big Numbers.Big Integers;
3
   use Ada.Numerics.Big Numbers.Big Integers;
4
5
   with Ada.Numerics.Big Numbers.Big Reals;
6
   use Ada.Numerics.Big Numbers.Big Reals;
7
8
   procedure Show Simple Big Numbers is
9
       BI : Big Integer;
10
       BR : Big Real;
11
   begin
12
       BI := 12345678901234567890;
13
       BR := 2.0 ** 1234;
14
15
       Put Line ("BI: " & BI'Image);
16
       Put_Line ("BR: " & BR'Image);
17
18
       BI := BI + 1;
19
       BR := BR + 1.0;
20
21
       Put_Line ("BI: " & BI'Image);
Put_Line ("BR: " & BR'Image);
22
23
   end Show Simple Big Numbers;
24
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Simple_Big_Numbers
MD5: b6a5e9ad170b09cbbabeb3ce06cc958c

Runtime output

```
BI: 12345678901234567890
BR:

-29581122460809862906004469571610359078633968713537299223955620705065735079623892426105383724837

-000
BI: 12345678901234567891
BR:

-29581122460809862906004469571610359078633968713537299223955620705065735079623892426105383724837

-000
```

In this example, we're declaring the big integer BI and the big real BR, and we're incrementing them by one.

Naturally, we're not limited to using the + operator (such as in this example). We can use the same operators on big numbers that we can use with normal numeric types. In fact, the common unary operators (+, -, abs) and binary operators (+, -, *, /, **, Min and Max) are available to us. For example:

Listing 77: show_simple_big_numbers_operators.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Numerics.Big_Numbers.Big_Integers;
duse Ada.Numerics.Big_Numbers.Big_Integers;
procedure Show_Simple_Big_Numbers_Operators is
```

```
BI : Big_Integer;
7
   begin
8
      BI := 12345678901234567890;
9
10
      Put_Line ("BI: " & BI'Image);
11
12
      BI := -BI + BI / 2;
13
      BI := BI - BI * 2;
14
15
      Put_Line ("BI: " & BI'Image);
16
   end Show_Simple_Big_Numbers_Operators;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Simple_Big_Numbers_

Goperators

MD5: 198708787bfcd6e16ec4fba718706af6
```

Runtime output

```
BI: 12345678901234567890
BI: 6172839450617283945
```

In this example, we're applying the four basic operators (+, -, *, /) on big integers.

8.7.2 Factorial

A typical example is the factorial¹⁶⁹: a sequence of the factorial of consecutive small numbers can quickly lead to big numbers. Let's take this implementation as an example:

Listing 78: factorial.ads

```
1 function Factorial (N : Integer)
2 return Long_Long_Integer;
```

Listing 79: factorial.adb

```
function Factorial (N : Integer)
1
                         return Long Long Integer is
2
      Fact : Long_Long_Integer := 1;
3
   begin
4
      for I in 2 .. N loop
5
         Fact := Fact * Long_Long_Integer (I);
6
      end loop;
7
8
      return Fact;
9
  end Factorial;
10
```

Listing 80: show_factorial.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Factorial;
procedure Show_Factorial is
begin
for I in 1 .. 50 loop
Put_Line (I'Image & "! = "
```

(continues on next page)

¹⁶⁹ https://en.wikipedia.org/wiki/Factorial

```
9 & Factorial (I)'Image);
10 end loop;
```

```
10 end Show_Factorial;
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Factorial_Integer
MD5: 9b20469533706ef025a03b506a07b920

Runtime output

1! = 1 2! = 2 3! = 6 4! = 24 5! = 120 6! = 720 7! = 5040 8! = 40320 9! = 362880 10! = 362880011! = 3991680012! = 47900160013! = 6227020800 14! = 87178291200 15! = 130767436800016! = 2092278988800017! = 35568742809600018! = 6402373705728000 19! = 12164510040883200020! = 2432902008176640000

raised CONSTRAINT_ERROR : factorial.adb:6 overflow check failed

Here, we're using Long_Long_Integer for the computation and return type of the Factorial function. (We're using Long_Long_Integer because its range is probably the biggest possible on the machine, although that is not necessarily so.) The last number we're able to calculate before getting an exception is 20!, which basically shows the limitation of standard integers for this kind of algorithm. If we use big integers instead, we can easily display all numbers up to 50! (and more!):

Listing 81: factorial.ads

```
with Ada.Numerics.Big_Numbers.Big_Integers;
use Ada.Numerics.Big_Numbers.Big_Integers;
function Factorial (N : Integer)
return Big_Integer;
```

Listing 82: factorial.adb

```
function Factorial (N : Integer)
1
                         return Big_Integer is
2
      Fact : Big_Integer := 1;
3
   begin
4
      for I in 2 .. N loop
5
          Fact := Fact * To_Big_Integer (I);
6
      end loop;
7
8
      return Fact;
9
  end Factorial;
10
```

```
Listing 83: show_big_number_factorial.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
  with Factorial;
3
4
  procedure Show_Big_Number_Factorial is
5
   begin
6
      for I in 1 .. 50 loop
7
         Put_Line (I'Image & "! = "
8
                   & Factorial (I)'Image);
9
      end loop;
10
  end Show_Big_Number_Factorial;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Factorial_Big_Numbers
MD5: d1f6464a3232d574d01f7ac14b822731

Runtime output

1! = 1
2! = 2
3! = 6
4! = 24
5! = 120
6! = 720
7! = 5040
8! = 40320
9! = 362880
10! = 3628800
11! = 39916800
12! = 479001600
13! = 6227020800
14! = 87178291200
15! = 1307674368000
16! = 20922789888000
17! = 355687428096000
18! = 6402373705728000
19! = 121645100408832000
20! = 2432902008176640000
21! = 51090942171709440000
22! = 1124000727777607680000
23! = 25852016738884976640000
24! = 620448401733239439360000
25! = 15511210043330985984000000
26! = 403291461126605635584000000
27! = 10888869450418352160768000000
28! = 304888344611713860501504000000
29! = 8841761993739701954543616000000
30! = 265252859812191058636308480000000
31! = 8222838654177922817725562880000000
32! = 2631308369336935301672180121600000000
33! = 8683317618811886495518194401280000000
34! = 295232799039604140847618609643520000000
35! = 10333147966386144929666651337523200000000
36! = 371993326789901217467999448150835200000000
37! = 13763753091226345046315979581580902400000000
38! = 523022617466601111760007224100074291200000000
39! = 20397882081197443358640281739902897356800000000
40! = 815915283247897734345611269596115894272000000000

41! =	33452526613163807108170062053440751665152000000000
42! =	1405006117752879898543142606244511569936384000000000
43! =	60415263063373835637355132068513997507264512000000000
44! =	2658271574788448768043625811014615890319638528000000000
45! =	119622220865480194561963161495657715064383733760000000000
46! =	5502622159812088949850305428800254892961651752960000000000
47! =	258623241511168180642964355153611979969197632389120000000000
48! =	12413915592536072670862289047373375038521486354677760000000000
49! =	608281864034267560872252163321295376887552831379210240000000000
50! =	30414093201713378043612608166064768844377641568960512000000000000

As we can see in this example, replacing the **Long_Long_Integer** type by the Big_Integer type fixes the problem (the runtime exception) that we had in the previous version. (Note that we're using the To_Big_Integer function to convert from **Integer** to Big_Integer: we discuss these conversions next.)

Note that there is a limit to the upper bounds for big integers. However, this limit isn't dependent on the hardware types — as it's the case for normal numeric types —, but rather compiler specific. In other words, the compiler can decide how much memory it wants to use to represent big integers.

8.7.3 Conversions

Most probably, we want to mix big numbers and *standard* numbers (i.e. integer and real numbers) in our application. In this section, we talk about the conversion between big numbers and standard types.

Validity

The package specifications of big numbers include subtypes that *ensure* that a actual value of a big number is valid:

Туре	Subtype for valid values
Big Integers	Valid_Big_Integer
Big Reals	Valid_Big_Real

These subtypes include a contract for this check. For example, this is the definition of the Valid_Big_Integer subtype:

Any operation on big numbers is actually performing this validity check (via a call to the Is_Valid function). For example, this is the addition operator for big integers:

As we can see, both the input values to the operator as well as the return value are expected to be valid — the Valid_Big_Integer subtype triggers this check, so to say. This approach ensures that an algorithm operating on big numbers won't be using invalid values.

Conversion functions

These are the most important functions to convert between big number and *standard* types:

Category	To big number	From big number
Big Integers	• To_Big_Integer	 To_Integer (Integer) From_Big_Integer (other integer types)
Big Reals	 To_Big_Real (floating- point types or fixed- point types) 	• From_Big_Real
	 To_Big_Real (Valid_Big_Integer) To_Real (Integer) 	 Numerator, Denomina- tor (Integer)

In the following sections, we discuss these functions in more detail.

Big integer to integer

We use the To_Big_Integer and To_Integer functions to convert back and forth between Big_Integer and Integer types:

Listing 84: show_simple_big_integer_conversion.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Numerics.Big Numbers.Big Integers;
3
   use Ada.Numerics.Big_Numbers.Big_Integers;
4
5
   procedure Show Simple Big Integer Conversion is
6
       BI : Big_Integer;
7
       I : Integer := 10000;
8
9
   begin
       BI := To_Big_Integer (I);
10
       Put_Line ("BI: " & BI'Image);
11
12
       I := To_Integer (BI + 1);
Put_Line ("I: " & I'Image);
13
14
   end Show_Simple_Big_Integer_Conversion;
15
```

Code block metadata

Runtime output

BI: 10000 I: 10001

In addition, we can use the generic Signed_Conversions and Unsigned_Conversions packages to convert between Big_Integer and any signed or unsigned integer types:

Listing 85: show_arbitrary_big_integer_conversion.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Numerics.Big_Numbers.Big_Integers;
3
   use Ada.Numerics.Big Numbers.Big Integers;
4
5
   procedure Show Arbitrary Big Integer Conversion is
6
7
      type Mod 32 Bit is mod 2 ** 32;
8
9
      package Long_Long_Integer_Conversions is new
10
        Signed_Conversions (Long_Long_Integer);
11
      use Long_Long_Integer_Conversions;
12
13
      package Mod 32 Bit Conversions is new
14
        Unsigned Conversions (Mod 32 Bit);
15
      use Mod_32_Bit_Conversions;
16
17
          : Big Integer;
      BI
18
      LLI : Long_Long_Integer := 10000;
19
      U_32 : Mod_32_Bit
                                := 2 ** 32 + 1;
20
21
   begin
22
      BI := To_Big_Integer (LLI);
23
      Put_Line ("BI: " & BI'Image);
24
25
      LLI := From_Big_Integer (BI + 1);
26
      Put Line ("LLI: " & LLI'Image);
27
28
      BI := To Big Integer (U 32);
29
      Put_Line ("BI: " & BI'Image);
30
31
      U_32 := From_Big_Integer (BI + 1);
32
      Put_Line ("U_32: " & U_32'Image);
33
34
   end Show_Arbitrary_Big_Integer_Conversion;
35
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Arbitrary_Big_

Gamma Gamma
```

Runtime output

BI:	10000
LLI:	10001
BI:	1
U_32:	2

In this examples, we declare the Long_Long_Integer_Conversions and the Mod_32_Bit_Conversions to be able to convert between big integers and the Long_Long_Integer and the Mod_32_Bit types, respectively.

Note that, when converting from big integer to integer, we used the To_Integer function, while, when using the instances of the generic packages, the function is named From_Big_Integer.

Big real to floating-point types

When converting between big real and floating-point types, we have to instantiate the generic Float_Conversions package:

Listing 86: show_big_real_floating_point_conversion.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Numerics.Big Numbers.Big Reals;
3
   use Ada.Numerics.Big_Numbers.Big_Reals;
4
5
   procedure Show Big Real Floating Point Conversion
6
   is
7
      type D10 is digits 10;
8
9
      package D10_Conversions is new
10
         Float Conversions (D10);
11
      use D10_Conversions;
12
13
      package Long_Float_Conversions is new
14
         Float Conversions (Long Float);
15
      use Long Float Conversions;
16
17
      BR : Big_Real;
18
      LF : Long_Float := 2.0;
19
      F10 : D10
                       := 1.999;
20
21
   begin
22
      BR := To_Big_Real (LF);
23
      Put_Line ("BR: " & BR'Image);
24
25
      LF := From_Big_Real (BR + 1.0);
26
      Put Line ("LF: " & LF'Image);
27
28
      BR := To_Big_Real (F10);
29
                       " & BR'Image);
      Put_Line ("BR:
30
31
      F10 := From_Big_Real (BR + 0.1);
32
      Put_Line ("F10: " & F10'Image);
33
34
   end Show_Big_Real_Floating_Point_Conversion;
35
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Big_Real_Floating_ ⊶Point_Conversion MD5: 4ccb570b964d11d215660f5929f2709c

Runtime output

In this example, we declare the D10_Conversions and the Long_Float_Conversions to be able to convert between big reals and the custom floating-point type D10 and the Long_Float type, respectively. To do that, we use the To_Big_Real and the From_Big_Real functions.

Big real to fixed-point types

When converting between big real and ordinary fixed-point types, we have to instantiate the generic Fixed_Conversions package:

Listing 87: show_big_real_fixed_point_conversion.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Numerics.Big Numbers.Big Reals;
3
   use Ada.Numerics.Big_Numbers.Big_Reals;
4
5
   procedure Show Big Real Fixed Point Conversion
6
   is
7
      D : constant := 2.0 ** (-31);
8
      type TQ31 is delta D range -1.0 .. 1.0 - D;
9
10
      package TQ31_Conversions is new
11
        Fixed Conversions (TQ31);
12
      use TQ31 Conversions;
13
14
      BR : Big Real;
15
      FQ31 : TQ31 := 0.25;
16
17
   begin
18
      BR := To_Big_Real (FQ31);
19
      Put_Line ("BR: " & BR'Image);
20
21
      FQ31 := From_Big_Real (BR * 2.0);
22
      Put_Line ("FQ31: " & FQ31'Image);
23
24
   end Show_Big_Real_Fixed_Point_Conversion;
25
```

Code block metadata

Runtime output

BR: 0.250 FQ31: 0.500000000

In this example, we declare the TQ31_Conversions to be able to convert between big reals and the custom fixed-point type TQ31 type. Again, we use the To_Big_Real and the From_Big_Real functions for the conversions.

Note that there's no direct way to convert between decimal fixed-point types and big real types. (Of course, you could perform this conversion indirectly by using a floating-point or an ordinary fixed-point type in between.)

Big reals to (big) integers

We can also convert between big reals and big integers (or standard integers):

Listing 88: show_big_real_big_integer_conversion.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Numerics.Big_Numbers.Big_Integers;
duse Ada.Numerics.Big_Numbers.Big_Integers;
```

```
5
   with Ada.Numerics.Big_Numbers.Big_Reals;
6
   use Ada.Numerics.Big_Numbers.Big_Reals;
7
   procedure Show_Big_Real_Big_Integer_Conversion
9
10
   is
      I : Integer;
11
      BI : Big_Integer;
12
      BR : Big_Real;
13
14
   begin
15
      I := 12345;
16
      BR := To_Real (I);
17
      Put_Line ("BR (from I): " & BR'Image);
18
19
      BI := 123456;
20
      BR := To_Big_Real (BI);
21
      Put_Line ("BR (from BI): " & BR'Image);
22
23
   end Show_Big_Real_Big_Integer_Conversion;
24
```

Code block metadata

Runtime output

```
BR (from I): 12345.000
BR (from BI): 123456.000
```

Here, we use the To_Real and the To_Big_Real and functions for the conversions.

String conversions

In addition to that, we can use string conversions:

```
Listing 89: show_big_number_string_conversion.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Numerics.Big_Numbers.Big_Integers;
3
   use Ada.Numerics.Big_Numbers.Big_Integers;
4
5
   with Ada.Numerics.Big_Numbers.Big_Reals;
6
   use Ada.Numerics.Big_Numbers.Big_Reals;
7
8
   procedure Show_Big_Number_String_Conversion
9
10
   is
      BI : Big_Integer;
11
      BR : Big_Real;
12
   beain
13
      BI := From_String ("12345678901234567890");
14
      BR := From_String ("12345678901234567890.0");
15
16
      Put_Line ("BI: "
17
                 & To_String (Arg
                                    => BI,
18
                               Width => 5,
19
                               Base => 2));
20
      Put Line ("BR: "
21
```

```
22 & To_String (Arg => BR,
23 Fore => 2,
24 Aft => 6,
25 Exp => 18));
26 end Show_Big_Number_String_Conversion;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Big_Number_String_

⊶Conversion

MD5: aalf19af04b0b901a086ac86151693a7
```

Runtime output

In this example, we use the From_String to convert a string to a big number. Note that the From_String function is actually called when converting a literal — because of the corresponding aspect for user-defined literals in the definitions of the Big_Integer and the Big_Real types.

For further reading...

Big numbers are implemented using *user-defined literals* (page 70), which we discussed previously. In fact, these are the corresponding type declarations:

```
Declaration from
- -
-- Ada.Numerics.Big Numbers.Big Integers;
type Big Integer is private
 with Integer Literal => From Universal Image,
      Put Image
                  => Put Image;
function From Universal Image
  (Ara : Strina)
  return Valid Big Integer
    renames From String;
-- Declaration from
-- Ada.Numerics.Big Numbers.Big Reals;
type Big Real is private
 with Real Literal => From Universal Image,
                 => Put Image;
      Put Image
function From Universal Image
  (Arg : String)
   return Valid Big Real
     renames From String;
As we can see in these declarations, the From String function renames the
From Universal Image function, which is being used for the user-defined literals.
```

Also, we call the To_String function to get a string for the big numbers. Naturally, using the To_String function instead of the Image attribute — as we did in previous examples — allows us to customize the format of the string that we display in the user message.

8.7.4 Other features of big integers

Now, let's look at two additional features of big integers:

- the natural and positive subtypes, and
- other available operators and functions.

Big positive and natural subtypes

Similar to integer types, big integers have the Big_Natural and Big_Positive subtypes to indicate natural and positive numbers. However, in contrast to the Natural and Positive subtypes, the Big_Natural and Big_Positive subtypes are defined via predicates rather than the simple ranges of normal (ordinary) numeric types:

```
subtype Natural is
  Integer range 0 .. Integer'Last;
subtype Positive is
  Integer range 1 .. Integer'Last;
subtype Big Natural is Big Integer
  with Dynamic_Predicate =>
         (if Is Valid (Big Natural)
            then Big Natural \geq = 0),
       Predicate_Failure =>
         (raise Constraint_Error);
subtype Big_Positive is Big_Integer
 with Dynamic_Predicate =>
         (if Is_Valid (Big_Positive)
            then Big_Positive > 0),
       Predicate Failure =>
         (raise Constraint Error);
```

Therefore, we cannot simply use attributes such as Big_Natural'First. However, we can use the subtypes to ensure that a big integer is in the expected (natural or positive) range:

Listing 90: show_big_positive_natural.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Ada.Numerics.Big Numbers.Big Integers;
3
   use Ada.Numerics.Big Numbers.Big Integers;
4
5
   procedure Show_Big_Positive_Natural is
6
      BI, D, N : Big Integer;
7
   begin
8
      D := 3;
9
      N := 2;
10
      BI := Big_Natural (D / Big_Positive (N));
11
12
      Put Line ("BI: " & BI'Image);
13
  end Show_Big_Positive_Natural;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Big_Positive_Natural
MD5: 844b41f001c9aed9cb99decb221d93fd
```

Runtime output

BI: 1

By using the Big_Natural and Big_Positive subtypes in the calculation above (in the assignment to BI), we ensure that we don't perform a division by zero, and that the result of the calculation is a natural number.

8.7.5 Other operators for big integers

We can use the **mod** and **rem** operators with big integers:

Listing 91: show_big_integer_rem_mod.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Numerics.Big Numbers.Big Integers;
3
   use Ada.Numerics.Big Numbers.Big Integers;
4
5
   procedure Show Big Integer Rem Mod is
6
      BI : Big Integer;
7
8
   begin
      BI := 145 \mod (-4);
9
      Put Line ("BI (mod): " & BI'Image);
10
11
      BI := 145 rem (-4);
12
      Put Line ("BI (rem): " & BI'Image);
13
   end Show_Big_Integer_Rem_Mod;
14
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Big_Integer_Rem_Mod MD5: 7347b617c51a3782921d997b3cfd5d37

Runtime output

BI (mod): -5 BI (rem): 1

In this example, we use the mod and rem operators in the assignments to BI.

Moreover, there's a Greatest_Common_Divisor function for big integers which, as the name suggests, calculates the greatest common divisor of two big integer values:

Listing 92: show big integer greatest common divisor.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Numerics.Big Numbers.Big Integers;
3
   use Ada.Numerics.Big_Numbers.Big_Integers;
4
5
   procedure Show Big Integer Greatest Common Divisor
6
7
   is
      BI : Big_Integer;
8
   begin
9
      BI := Greatest_Common_Divisor (145, 25);
10
      Put Line ("BI: " & BI'Image);
11
12
   end Show Big Integer Greatest Common Divisor;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Big_Integer_Greatest_ GCommon_Divisor MD5: 27e2f7b4cbe20ec979b672f3e7edfdb7

Runtime output

BI: 5

In this example, we retrieve the greatest common divisor of 145 and 25 (i.e.: 5).

8.7.6 Big real and quotients

An interesting feature of big reals is that they support quotients. In fact, we can simply assign 2/3 to a big real variable. (Note that we're able to omit the decimal points, as we write 2/3 instead of 2.0 / 3.0.) For example:

Listing 93: show_big_real_quotient_conversion.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Numerics.Big Numbers.Big Reals;
3
   use Ada.Numerics.Big Numbers.Big Reals;
4
5
   procedure Show_Big_Real_Quotient_Conversion
6
7
   is
      BR
            : Big_Real;
8
   begin
9
      BR := 2 / 3;
10
          Same as:
11
       - -
       -- BR := From_Quotient_String ("2 / 3");
12
13
      Put Line ("BR:
                         " & BR'Image);
14
15
      Put Line ("Q:
16
                 & To Quotient String (BR));
17
18
      Put Line ("Q numerator:
19
                 & Numerator (BR) 'Image);
20
      Put Line ("Q denominator:
21
                 & Denominator (BR) 'Image);
22
   end Show Big Real Quotient Conversion;
23
```

Code block metadata

Runtime output

BR: 0.666 Q: 2 / 3 Q numerator: 2 Q denominator: 3

In this example, we assign 2 / 3 to BR — we could have used the From_Quotient_String function as well. Also, we use the To_Quotient_String to get a string that represents the quotient. Finally, we use the Numerator and Denominator functions to retrieve the values, respectively, of the numerator and denominator of the quotient (as big integers) of the big real variable.

8.7.7 Range checks

Previously, we've talked about the Big_Natural and Big_Positive subtypes. In addition to those subtypes, we have the In_Range function for big numbers:

```
Listing 94: show_big_numbers_in_range.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Numerics.Big_Numbers.Big_Integers;
3
   use Ada.Numerics.Big_Numbers.Big_Integers;
4
5
   with Ada.Numerics.Big Numbers.Big Reals;
6
        Ada.Numerics.Big_Numbers.Big_Reals;
   use
7
8
   procedure Show_Big_Numbers_In_Range is
9
10
      BI : Big_Integer;
11
      BR : Big_Real;
12
13
      BI From : constant Big Integer := 0;
14
              : constant Big_Integer := 1024;
      BI To
15
16
      BR_From : constant Big_Real := 0.0;
17
      BR To
              : constant Big_Real := 1024.0;
18
19
   begin
20
      BI := 1023;
21
      BR := 1023.9;
22
23
       if In_Range (BI, BI_From, BI_To) then
24
          Put_Line ("BI ("
25
                     & BI'Image
26
                     & ") is in the "
27
                     & BI_From'Image
28
                     & "
29
                     & BI_To'Image
30
                     & " range");
31
      end if;
32
33
      if In_Range (BR, BR_From, BR_To) then
34
          Put_Line ("BR ("
35
                     & BR'Image
36
                     & ") is in the "
37
                     & BR_From'Image
38
                     & " . . "
39
                     & BR_To'Image
40
                     & " range");
41
      end if;
42
43
   end Show_Big_Numbers_In_Range;
44
```

Code block metadata

Project: Courses.Advanced_Ada.Data_Types.Numerics.Big_Numbers.Big_Numbers_In_Range
MD5: ded52ef7e9ef13a83264940ff9d8bcb3

Runtime output

```
BI (1023) is in the 0 .. 1024 range
BR (1023.900) is in the 0.000 .. 1024.000 range
```

In this example, we call the In_Range function to check whether the big integer number

(BI) and the big real number (BR) are in the range between 0 and 1024.

Part II Control Flow

EXPRESSIONS

9.1 Expressions: Definition

According to the Ada Reference Manual, an expression "is a formula that defines the computation or retrieval of a value." Also, when an expression is evaluated, the computed or retrieved value always has an associated type known at compile-time.

Even though the definition above is very simple, Ada expressions are actually very flexible — and they can also be very complex. In fact, if you read the corresponding section¹⁷⁰ of the Ada Reference Manual, you'll quickly discover that they include elements such as relations, membership choices, terms and primaries. Some of these are classic elements of expressions in programming languages, although some of their forms are unique to Ada. In this section, we present examples of just some of these elements. For a complete overview, please refer to the Reference Manual.

In the Ada Reference Manual

```
    4.4 Expressions<sup>171</sup>
```

9.1.1 Relations and simple expressions

Expressions usually consist of relations, which in turn consist of simple expressions. (There are more details to this, but we'll keep it simple for the moment.) Let's see a code example with a few expressions, which we dissect into the corresponding grammatical elements — we're going to discuss them later:

Listing 1:	show	expression	elements.adb

```
procedure Show_Expression_Elements is
1
      type Mode is (Off, A, B, C, D);
2
3
      pragma Unreferenced (B, C, D);
4
5
      subtype Active Mode is Mode
6
         range Mode'Succ (Off) .. Mode'Last;
7
8
      M1, M2 : Mode;
9
      Dummy
                 : Boolean;
10
   begin
11
      M1 := A;
12
13
      Dummy :=
14
           M1 in Active Mode
15
```

(continues on next page)

¹⁷⁰ http://www.ada-auth.org/standards/22rm/html/RM-4-4.html
 ¹⁷¹ http://www.ada-auth.org/standards/22rm/html/RM-4-4.html

```
and then M2 in Off | A;
16
17
      - -
           ^^^^ relation
      - -
18
      - -
19
                            ^^^^ relation
      - -
20
          ^^^^
      - -
21
22
      - -
                                        expression
23
      Dummy :=
24
         M1 in Active_Mode;
25
         ^^ name
26
      -- ^^ primary
27
         ^^ factor
28
      - -
         ^^ term
29
      - -
         ^^ simple expression
30
      - -
31
      - -
               ^^^^ membership choice
      - -
32
               ^^^^ membership choice list
      - -
33
34
      - -
         relation
      - -
35
         expression
      - -
36
37
      Dummy :=
38
         M2 in Off | A;
39
          ^^ name
40
         ^^ primary
      - -
41
         ^^ factor
      - -
42
         ^^ term
43
      - -
         ^^ simple expression
44
      - -
45
      - -
               ^^^ membership choice
      - -
46
                    ^ membership choice
      - -
47
               ^^^^ membership choice list
      - -
48
49
         ^^^^ relation
      - -
50
         expression
      - -
51
52
   end Show_Expression_Elements;
53
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Expressions_Definition.

←Expression_Elements

MD5: a22e6f2d2bc181ce77097a1de204eb62
```

Build output

In this code example, we see three expressions. As we mentioned earlier, every expression has a type; here, the type of each expression is **Boolean**.

The first expression (M1 in Active_Mode and then M2 in Off | A) consists of two relations: M1 in Active_Mode and M2 in Off | A. Let's discuss some of the details.

The M1 in Active_Mode relation consists of the simple expression M1 and the membership choice list Active_Mode. (Here, the in keyword is part of the relation definition.) Also, as we see in the comments of the source code, the simple expression M1 is, at the same time, a term, a factor, a primary and a name.

Let's briefly talk about this chain of syntactic elements for simple expressions. Very roughly

said, this is how we can break up simple expressions:

- a simple expression consists of terms;
- a term consists of factors;
- a factor consists of primaries;
- a primary can be one of those:
 - a numeric literal;
 - null;
 - a string literal;
 - an aggregate (page 247);
 - a name;
 - an allocator (like new Integer);
 - a parenthesized expression (page 431);
 - a conditional expression (page 433);
 - a quantified expression (page 436);
 - a declare expression (page 440).

For further reading...

The definition of simple expressions we've just seen is very simplified. In actuality, these are the grammatical elements specified in the Ada Reference Manual:

```
simple_expression ::=
  [unary_adding_operator] term {binary_adding_operator term}
term ::= factor {multiplying_operator factor}
factor ::= primary [** primary] | abs primary | not primary
primary ::=
  numeric_literal | null | string_literal | aggregate
| name | allocator | (expression)
| (conditional_expression) | (quantified_expression)
| (declare_expression)
```

Later on in this chapter, we discuss *conditional expressions* (page 433), *quantified expressions* (page 436) and *declare expressions* (page 440) in more details.

In the relation M2 in Off \mid A from the code example, Off \mid A is a membership choice list, and Off and A are membership choices.

1 For further reading...

Relations can actually be much more complicated than the one we just saw. In fact, this is the definition from the Ada Reference Manual: expression ::=

relation {and relation}
| relation {and then relation}
| relation {or relation}
| relation {or else relation}
| relation {xor relation}

```
relation ::=
    simple_expression
    [relational_operator simple_expression]
    simple_expression [not] in
    membership_choice_list
    | raise_expression
```

Again, for more details, please refer to the section on expressions¹⁷² of the Ada Reference Manual.

In the Ada Reference Manual

- 4.4 Expressions¹⁷³
- 4.5.2 Relational Operators and Membership Tests¹⁷⁴

9.1.2 Numeric expressions

The expressions we've seen so far had the **Boolean** type. Although much of the grammar described in the Manual exists exclusively for Boolean operations, we can also write numeric expressions such as the following one:

Listing 2:	show	numeric	expressions.adk	С
	_		_ '	

```
procedure Show_Numeric_Expressions is
1
      C1 : constant Integer := 5;
2
      Dummy :
                Integer;
3
   begin
4
      Dummy :=
5
       -2 ** 4 + 3 * C1 ** 8;
6
                            ^ numeric literal
      - -
7
                             ^ primary
      - -
8
                      name
      - -
9
                       ^^ primary
^^^ factor
      - -
10
11
      - -
                     ^ multiplying operator
12
      - -
                   ~
                        numeric literal
13
      - -
                   ^
                             primary
      - -
14
                   ~
                               factor
      - -
15
                   ..... term
      - -
16
      - -
17
            ^ numeric literal
^ primary
      - -
18
      - -
19
      -- ^ numeric literal
20
      -- ^ primary
21
          _____
      - -
                               factor
22
          ~~~~~
      - -
                               term
23
          ^ binary adding operator
      - -
24
         ^ unary adding operator
25
      - -
26
      - -
         simple expression
      - -
27
28
      - -
          expression
29
   end Show Numeric Expressions;
30
```

Code block metadata

¹⁷² http://www.ada-auth.org/standards/22rm/html/RM-4-4.html

¹⁷³ http://www.ada-auth.org/standards/22rm/html/RM-4-4.html

¹⁷⁴ http://www.ada-auth.org/standards/22rm/html/RM-4-5-2.html

Project: Courses.Advanced_Ada.Control_Flow.Expressions.Expressions_Definition. ⇔Numeric_Expressions MD5: a3c902c7aa5b0afe30ae220256c3306a

In this code example, the expression - 2 ** 4 + 3 * C1 ** 8 consists of just a single simple expression. (Note that simple expressions do not have to be "simple".) This simple expression consists of two terms: 2 ** 4 and 3 * C1 ** 8. While the 2 ** 4 term is also a single factor, the 3 * C1 ** 8 term consists of two factors: 3 and C1 ** 8. Both the 2 ** 4 and the C1 ** 8 factors consists of two primaries each:

- the 2 ** 4 factor has the primaries 2 and 4,
- the C1 ** 8 factor has the primaries C1 and 8.

1 In the Ada Reference Manual

4.4 Expressions¹⁷⁵

9.1.3 Other expressions

Expressions aren't limited to the **Boolean** type or to numeric types. Indeed, expressions can be of any type, and the definition of primaries we've seen earlier on already hints in this direction — as it includes elements such as allocators. Because expressions are very flexible, covering all possible variations and combinations in this section is out of scope. Again, please refer to the section on expressions¹⁷⁶ of the Ada Reference Manual for further details.

9.1.4 Parenthesized expression

An interesting aspect of primaries is that, by using parentheses, we can embed an expression inside another expression. As an example, let's discuss the following expression and its elements:

Listing 3: show_parenthesized_expressions.adb

```
procedure Show Parenthesized Expressions is
1
       C1 : constant Integer := 4;
2
       C2 : constant Integer := 5;
3
4
5
       Dummy : Integer;
6
   begin
       Dummy :=
7
            (2 + C1) * C2;
8
                  ~~
9
                           name
                 ~~
       - -
                            primary
10
                 ~~
                            factor
       - -
11
                 ~~
       - -
                            term
12
       - -
13
             ^
                            numeric literal
       - -
14
             ^
       - -
                            primary
15
             ~
16
       - -
                            factor
             ~
17
       - -
                            term
18
       - -
               ~
                            binary adding operator
19
       - -
            ~~~~~
                            simple expression
       - -
20
21
```

(continues on next page)

¹⁷⁵ http://www.ada-auth.org/standards/22rm/html/RM-4-4.html
 ¹⁷⁶ http://www.ada-auth.org/standards/22rm/html/RM-4-4.html

```
^^^^^
      - -
                     expression
22
                  primary
         ^^^^^
23
      - -
         ^^^^^
      - -
                      factor
24
      - -
25
                   ^^ factor
      - -
26
         ..... term
      - -
27
28
      - -
         ^^^^^ simple expression
      - -
29
      - -
30
         ••••••• expression
31
  end Show Parenthesized Expressions;
32
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Expressions.Expressions_Definition. ⇔Parenthesized_Expressions MD5: 5871d2b0cd33e4f562b96381e0f0d293

In this example, we first start with the single expression (2 + C1) * C2, which is also a simple expression consisting of just one term, which consists of two factors: (2 + C1) and C2. The (2 + C1) factor is also a primary. Now, because of the parentheses, we identify that the primary (2 + C1) is an expression that is embedded in another expression.

1 Important

To be fair, the existence of parentheses in a primary could also indicate other kinds of expressions, such as conditional or quantified expressions. However, differentiating between them is straightforward, as we'll see later on in this chapter.

We then proceed to parse the (2 + C1) expression, which consists of the terms 2 and C1. As we've seen in the comments of the code example, each of these terms consists of one factor, which consists of one primary. In the end, after parsing the primaries, we identify that 2 is a numeric literal and C1 is a name.

Note that the usage of parentheses might lead to situations where we have expressions in potentially unsuspected places. For example, consider the following code example:

Listing 4: show_name_in_expression.adb

```
procedure Show Name In Expression is
1
       type Mode is (Off, A, B, C, D);
2
3
      M1 : Mode;
4
   begin
5
      M1 := A;
6
7
      case M1 is
8
        when Off | D =>
9
           null;
10
        when A | B | C =>
11
          M1 := D;
12
      end case;
13
14
   end Show Name In Expression;
15
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Expressions.Expressions_Definition.Name_ □In_Expression

```
MD5: ec8fcbc511e6a372da4f0ad99d2619a5
```

Here, the case statement expects a selecting expression. In this case, M1 is identified as a name — after being identified as a relation, a simple expression, a term, a factor and a primary.

However, if we replace **case** M1 **is** by **case** (M1) **is**, (M1) is identified as a parenthesized expression, not as a name! This parenthesized expression is first parsed and evaluated, which might have implications in case statements, as we'll see *in another chapter* (page 458).

Let's look at another example, this time with a subprogram call:

Listing 5: increment_by_one.ads

```
procedure Increment_By_One (I : in out Integer);
```

Listing 6: increment_by_one.adb

```
procedure Increment_By_One (I : in out Integer) is
begin
I := I + 1;
end Increment_By_One;
```

Listing 7: show_name_in_expression.adb

```
with Increment_By_One;
procedure Show_Name_In_Expression is
V : Integer := 0;
begin
Increment_By_One ((V));
end Show_Name_In_Expression;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Expressions_Definition.Name_

⊲In_Expression

MD5: 4805df49dc702e5cb365252e58742dd2
```

Build output

```
show_name_in_expression.adb:6:23: error: actual for "I" must be a variable
gprbuild: *** compilation phase failed
```

The Increment_By_One procedure from this example expects a variable as an actual parameter because the parameter mode is **in out**. However, the (V) in the call to the procedure is interpreted as an expression, so we end up providing a value — the result of the expression — as the actual parameter instead of the V variable. Naturally, this is a compilation error. (Of course, writing Increment_By_One (V) fixes the error.)

9.2 Conditional Expressions

As we've seen before, we can write simple expressions such as I = 0 or D.Valid. A conditional expression, as the name implies, is an expression that contains a condition. This might be an "if-expression" (in the **if** ... **then** ... **else** form) or a "case-expression" (in the **case** ... **is when** => form).

The Max function in the following code example is an expression function implemented with a conditional expression — an if-expression, to be more precise:

```
Listing 8: expr_func.ads
```

```
package Expr_Func is
function Max (A, B : Integer) return Integer is
(if A >= B then A else B);
end Expr_Func;
```

Let's say we have a system with four states Off, On, Waiting, and Invalid. For this system, we want to implement a function named Toggled that returns the *toggled* value of a state S. If the current value of S is either Off or On, the function toggles from Off to On (or from On to Off). For other values, the state remains unchanged — i.e. the returned value is the same as the input value. This is the implementation using a conditional expression:

Listing 9: expr_func.ads

```
package Expr_Func is
1
2
      type State is (Off, On, Waiting, Invalid);
3
4
      function Toggled (S : State) return State is
5
         (if S = Off
6
           then On
7
           elsif S = On
8
             then Off
9
             else S);
10
11
   end Expr_Func;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Conditional_Expressions.

Gonditional_If_Expressions_1

MD5: 7a99711afecc0b481557f9874dfbf4de
```

As you can see, if-expressions may contain an **elsif** branch (and therefore be more complicated).

The code above corresponds to this more verbose version:

Listing 10: expr_func.ads

```
package Expr_Func is
type State is (Off, On, Waiting, Invalid);
function Toggled (S : State) return State;
end Expr_Func;
```

Listing 11: expr_func.adb

```
package body Expr_Func is
function Toggled (S : State) return State is
begin
find if S = Off then
foreturn On;
```

```
7 elsif S = On then
8 return Off;
9 else
10 return S;
11 end if;
12 end Toggled;
13
14 end Expr_Func;
```

Code block metadata

If we compare the if-block of this code example to the if-expression of the previous example, we notice that the if-expression is just a simplified version without the **return** keyword and the **end if**;. In fact, converting an if-block to an if-expression is quite straightforward.

We could also replace the if-expression used in the Toggled function above with a case-expression. For example:

Listing 12: expr_func.ads

```
package Expr_Func is
1
2
      type State is (Off, On, Waiting, Invalid);
3
4
      function Toggled (S : State) return State is
5
         (case S is
6
           when Off
                        => On,
7
           when On
                       => Off,
8
           when others => S);
9
10
   end Expr_Func;
11
```

Code block metadata

Note that we use commas in case-expressions to separate the alternatives (the **when** expressions). The code above corresponds to this more verbose version:

Listing 13: expr_func.ads

```
package Expr_Func is
type State is (Off, On, Waiting, Invalid);
function Toggled (S : State) return State;
end Expr_Func;
```

Listing 14: expr_func.adb

```
package body Expr_Func is
function Toggled (S : State) return State is
begin
```

```
case S is
5
            when Off
                      => return On;
6
            when On
                       => return Off;
7
            when others => return S;
8
         end case;
9
      end Toggled;
10
11
  end Expr_Func;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Conditional_Expressions.

→Conditional_Case_Expressions_2

MD5: db6a0737e3931c83c31f53e4da3d8a2b
```

If we compare the case block of this code example to the case-expression of the previous example, we notice that the case-expression is just a simplified version of the case block without the **return** keyword and the **end case;**, and with alternatives separated by commas instead of semicolons.

1 In the Ada	Reference Manual
• 4.5.7 Con	ditional Expressions ¹⁷⁷

9.3 Quantified Expressions

Quantified expressions are **for** expressions using a quantifier — which can be either **all** or **some** — and a predicate. This kind of expressions let us formalize statements such as:

- "all values of array A must be zero" into for all I in A'Range => A (I) = 0, and
- "at least one value of array A must be zero" into for some I in A'Range => A (I)
 = 0.

In the quantified expression **for all** I **in** A'Range => A (I) = 0, the quantifier is **all** and the predicate is A (I) = 0. In the second expression, the quantifier is **some**. The result of a quantified expression is always a Boolean value.

For example, we could use the quantified expressions above and implement these two functions:

- Is_Zero, which checks whether all components of an array A are zero, and
- Has_Zero, which checks whether array A has at least one component of the array A is zero.

This is the complete code:

Listing 15: int_arrays.ads

```
package Int_Arrays is
type Integer_Arr is
array (Positive range <>) of Integer;
function Is_Zero (A : Integer_Arr)
return Boolean is
(for all I in A'Range => A (I) = 0);
```

(continues on next page)

¹⁷⁷ http://www.ada-auth.org/standards/22rm/html/RM-4-5-7.html

Listing 16: int_arrays.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Int Arrays is
3
4
      procedure Display Array (A
                                    : Integer Arr;
5
                                 Name : String) is
6
      begin
7
          Put (Name & ": ");
8
          for E of A loop
9
             Put (E'Image & " ");
10
          end loop;
11
          New_Line;
12
      end Display_Array;
13
14
   end Int_Arrays;
15
```

Listing 17: test_int_arrays.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Int_Arrays; use Int_Arrays;
3
4
   procedure Test_Int_Arrays is
5
      A : Integer_Arr := (0, 0, 1);
6
   begin
7
      Display_Array (A, "A");
8
      Put_Line ("Is_Zero: "
9
                 & Boolean'Image (Is_Zero (A)));
10
      Put_Line ("Has_Zero: "
11
                 & Boolean'Image (Has_Zero (A)));
12
13
      A := (0, 0, 0);
14
15
      Display_Array (A, "A");
16
      Put_Line ("Is_Zero: "
17
                 & Boolean'Image (Is_Zero (A)));
18
      Put_Line ("Has_Zero: "
19
                 & Boolean'Image (Has_Zero (A)));
20
   end Test_Int_Arrays;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Quantified_Expression.

→Quantified_Expression_1

MD5: 4bbda8a3830272748500f797f23f76fc
```

Runtime output

A: 0 0 1 Is_Zero: FALSE Has_Zero: TRUE A: 0 0 0 Is_Zero: TRUE Has_Zero: TRUE

As you might have expected, we can rewrite a quantified expression as a loop in the **for** I **in** A'Range **loop if** ... **return** ... form. In the code below, we're implementing Is_Zero and Has_Zero using loops and conditions instead of quantified expressions:

Listing 18: int_arrays.ads

```
package Int_Arrays is
1
2
      type Integer_Arr is
3
        array (Positive range <>) of Integer;
4
5
      function Is_Zero (A : Integer_Arr)
6
                          return Boolean;
7
8
      function Has_Zero (A : Integer_Arr)
9
                           return Boolean;
10
11
      procedure Display_Array (A : Integer_Arr;
12
                                 Name : String);
13
14
   end Int_Arrays;
15
```

Listing 19: int arrays.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Int_Arrays is
3
4
5
       function Is_Zero (A : Integer_Arr)
                          return Boolean is
6
       begin
7
          for I in A'Range loop
8
             if A (I) /= 0 then
9
                return False;
10
             end if;
11
          end loop;
12
13
          return True;
14
       end Is_Zero;
15
16
       function Has_Zero (A : Integer_Arr)
17
                           return Boolean is
18
       begin
19
          for I in A'Range loop
20
            if A(I) = 0 then
21
               return True;
22
            end if;
23
          end loop;
24
25
          return False;
26
       end Has_Zero;
27
28
       procedure Display_Array (A : Integer_Arr;
29
                                  Name : String) is
30
```

```
begin
31
          Put (Name & ": ");
32
          for E of A loop
33
             Put (E'Image & " ");
34
          end loop;
35
          New_Line;
36
       end Display_Array;
37
38
   end Int_Arrays;
39
```

Listing 20: test_int_arrays.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Int_Arrays; use Int_Arrays;
3
4
   procedure Test Int Arrays is
5
      A : Integer Arr := (0, 0, 1);
6
   begin
7
      Display_Array (A, "A");
8
      Put_Line ("Is_Zero: '
9
                 & Boolean'Image (Is_Zero (A)));
10
      Put Line ("Has Zero: "
11
                 & Boolean 'Image (Has_Zero (A)));
12
13
      A := (0, 0, 0);
14
15
      Display_Array (A, "A");
16
      Put_Line ("Is_Zero: "
17
                 & Boolean'Image (Is_Zero (A)));
18
      Put_Line ("Has_Zero: "
19
                 & Boolean 'Image (Has_Zero (A)));
20
   end Test_Int_Arrays;
21
```

Code block metadata

Runtime output

A: 0 0 1 Is_Zero: FALSE Has_Zero: TRUE A: 0 0 0 Is_Zero: TRUE Has_Zero: TRUE

So far, we've seen quantified expressions using indices — e.g. for all I in A'Range => We could avoid indices in quantified expressions by simply using the E of A form. In this case, we can just write for all E of A => Let's adapt the implementation of Is_Zero and Has_Zero using this form:

Listing 21: int_arrays.ads

```
1 package Int_Arrays is
2
3 type Integer_Arr is
4 array (Positive range <>) of Integer;
```

```
5
       function Is_Zero (A : Integer_Arr)
6
                           return Boolean is
7
          (for all E of A \Rightarrow E = 0);
8
9
       function Has_Zero (A : Integer_Arr)
10
                            return Boolean is
11
          (for some E of A => E = 0);
12
13
   end Int_Arrays;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Quantified_Expression.

⊲Quantified_Expression_3

MD5: 059d12a6529483ebcc5db23dc6262896
```

Here, we're checking the components E of the array A and comparing them against zero.

1 In the Ada Reference Manual

4.5.8 Quantified Expressions¹⁷⁸

9.4 Declare Expressions

So far, we've seen expressions that make use of existing objects declared outside of the expression. Sometimes, we might want to declare constant objects inside the expression, so we can use them locally in the expression. Similarly, we might want to rename an object and use the renamed object in an expression. In those cases, we can use a declare expression.

A declare expression allows for declaring or renaming objects within an expression:

Listing 22: p.ads

```
package P is
function Max (A, B : Integer) return Integer is
(declare
Bigger_A : constant Boolean := (A >= B);
begin
(if Bigger_A then A else B));
```

end P;

Code block metadata

The declare expression starts with the **declare** keyword and the usual object declarations, and it's followed by the **begin** keyword and the body. In this example, the body of the declare expression is a conditional expression.

Of course, the code above isn't really useful, so let's look at a more complete example:

¹⁷⁸ http://www.ada-auth.org/standards/22rm/html/RM-4-5-8.html

```
Listing 23: integer_arrays.ads
```

```
package Integer_Arrays is
1
2
       type Integer_Array is
3
         array (Positive range <>) of Integer;
4
5
       function Sum (Arr : Integer Array)
6
                      return Integer;
7
8
9
       -- Expression function using
10
          declare expression:
       - -
11
12
       function Avg (Arr : Integer_Array)
13
                      return Float is
14
         (declare
15
             A :
                           Integer Array renames Arr;
16
             S : constant Float := Float (Sum (A));
17
             L : constant Float := Float (A'Length);
18
          begin
19
             S / L);
20
21
   end Integer_Arrays;
22
```

Listing 24: integer_arrays.adb

```
package body Integer_Arrays is
1
2
      function Sum (Arr : Integer_Array)
3
                      return Integer is
4
      begin
5
         return Acc : Integer := 0 do
6
             for V of Arr loop
7
                Acc := Acc + V;
8
             end loop;
9
         end return;
10
      end Sum;
11
12
   end Integer_Arrays;
13
```

Listing 25: show_integer_arrays.adb

```
with Ada.Text_I0;
                          use Ada.Text_I0;
1
2
   with Integer_Arrays; use Integer_Arrays;
3
4
   procedure Show_Integer_Arrays is
5
      Arr : constant Integer_Array := [1, 2, 3];
6
   begin
7
      Put Line ("Sum: "
8
                 & Sum (Arr) 'Image);
9
      Put Line ("Avg:
10
                 & Avg (Arr)'Image);
11
   end Show_Integer_Arrays;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Declare_Expressions.Integer_

⇔Arrays

MD5: 30a035038508549822c819b60638133d
```

Runtime output

Sum: 6 Avg: 2.00000E+00

In this example, the Avg function is implemented using a declare expression. In this expression, A renames the Arr array, and S is a constant initialized with the value returned by the Sum function.

In the Ada Reference Manual

```
    4.5.9 Declare Expressions<sup>179</sup>
```

9.4.1 Restrictions in the declarative part

The declarative part of a declare expression is more restricted than the declarative part of a subprogram or declare block. In fact, we cannot:

- declare variables;
- · declare constants of limited types;
- · rename an object of limited type that is constructed within the declarative part;
- declare aliased constants;
- declare constants that make use of the Access or Unchecked_Access attributes in the initialization;
- declare constants of anonymous access type.

Let's see some examples of erroneous declarations:

```
Listing 26: integer_arrays.ads
```

```
package Integer_Arrays is
1
2
      type Integer Array is
3
        array (Positive range <>) of Integer;
4
5
      type Integer_Sum is limited private;
6
7
      type Const_Integer_Access is
8
         access constant Integer;
9
10
      function Sum (Arr : Integer_Array)
11
                      return Integer;
12
13
      function Sum (Arr : Integer_Array)
14
                      return Integer_Sum;
15
16
17
          Expression function using
       - -
18
          declare expression:
       - -
19
20
      function Avg (Arr : Integer_Array)
21
                      return Float is
22
         (declare
23
             A : Integer_Array renames Arr;
24
25
             S1 : aliased constant Integer := Sum (A);
26
```

(continues on next page)

¹⁷⁹ http://www.ada-auth.org/standards/22rm/html/RM-4-5-9.html

```
ERROR: aliased constant
27
              - -
28
             S : Float := Float (S1);
29
             L : Float := Float (A'Length);
30
              -- ERROR: declaring variables
31
32
             S2 : constant Integer_Sum := Sum (A);
33
              -- ERROR: declaring constant of
34
                          limited type
35
36
             A1 : Const_Integer_Access :=
37
                     S1'Unchecked_Access;
38
                  ERROR: using 'Unchecked_Access
              - -
39
                         attribute
40
              - -
41
             A2 : access Integer := null;
42
             -- ERROR: declaring object of
43
                         anonymous access type
44
          begin
45
             S / L);
46
47
   private
48
49
       type Integer_Sum is new Integer;
50
51
   end Integer_Arrays;
52
```

Listing 27: integer arrays.adb

```
package body Integer_Arrays is
1
2
      function Sum (Arr : Integer_Array)
3
                      return Integer is
4
      begin
5
          return Acc : Integer := 0 do
6
             for V of Arr loop
7
                Acc := Acc + V;
8
             end loop;
9
          end return;
10
      end Sum;
11
12
      function Sum (Arr : Integer_Array)
13
                      return Integer_Sum is
14
         (Integer_Sum (Integer'(Sum (Arr))));
15
16
   end Integer Arrays;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Declare_Expressions.Integer_
_Arrays_Error
MD5: ea38f5067c849b85685d70ffc386f7a7
```

Build output

(continued from previous page) integer_arrays.ads:42:15: error: anonymous access type not allowed in declare_ ⇔expression gprbuild: *** compilation phase failed

In this version of the Avg function, we see many errors in the declarative part of the declare expression. If we convert the declare expression into an actual function implementation, however, those declarations won't trigger compilation errors. (Feel free to try this out!)

9.5 Reduction Expressions

\rm 1 Note

This feature was introduced in Ada 2022.

A reduction expression reduces a list of values into a single value. For example, we can reduce the list [2, 3, 4] to a single value:

- by adding the values of the list: 2 + 3 + 4 = 9, or
- by multiplying the values of the list: 2 * 3 * 4 = 24.

We write a reduction expression by using the Reduce attribute and providing the reducer and its initial value:

- the reducer is the operator (e.g.: + or *) that we use to *combine* the values of the list;
- the initial value is the value that we use before all other values of the list.

For example, if we use + as the operator and 0 an the initial value, we get the reduction expression: 0 + 2 + 3 + 4 = 9. This can be implemented using an array:

Listing 28: show_reduction_expression.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Reduction_Expression is
3
      A : array (1 .. 3) of Integer;
4
      I : Integer;
5
   begin
6
      A := [2, 3, 4];
7
      I := A'Reduce ("+", 0);
8
9
      Put Line ("A = "
10
                 & A'Image);
11
      Put_Line ("I = "
12
                 & I'Image);
13
   end Show_Reduction_Expression;
14
```

Code block metadata

Runtime output

A = [2, 3, 4] I = 9 Here, we have the array A with a list of values. The A'Reduce ("+", 0) expression reduces the list of values of A into a single value — in this case, an integer value that is stored in I. This statement is equivalent to:

```
I := 0;
for E of A loop
I := I + E;
end loop;
```

Naturally, we can reduce the array using the * operator:

Listing 29: show_reduction_expression.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Reduction_Expression is
3
      A : array (1 .. 3) of Integer;
4
      I : Integer;
5
   begin
6
7
      A := [2, 3, 4];
      I := A'Reduce ("*", 1);
8
9
      Put_Line ("A = "
10
                 & A'Image);
11
      Put Line ("I = "
12
                 & I'Image);
13
   end Show_Reduction_Expression;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Reduction_Expressions.

→Simple_Reduction_Expression

MD5: 98elde10863eed4bd12cc6ab1d7ce7ef
```

Runtime output

A = [2, 3, 4] I = 24

In this example, we call A'Reduce ("*", 1) to reduce the list. (Note that we use an initial value of one because it is the identity element¹⁸⁰ of a multiplication, so the complete operation is: 1 * 2 * 3 * 4 = 24.)

In the Ada Reference Manual

Reduction Expressions¹⁸¹

9.5.1 Value sequences

In addition to arrays, we can apply reduction expression to value sequences, which consist of an iterated element association — for example, [for I in 1 ... 3 => I + 1]. We can simply *append* the reduction expression to a value sequence:

¹⁸⁰ https://en.wikipedia.org/wiki/Identity_element

¹⁸¹ http://www.ada-auth.org/standards/22rm/html/RM-4-5-10.html

```
Listing 30: show_reduction_expression.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Reduction Expression is
3
      I : Integer;
4
   begin
5
      I := [for I in 1 .. 3 =>
6
               I + 1]'Reduce ("+", 0);
7
      Put Line ("I = "
8
                 & I'Image);
9
10
      I := [for I in 1 .. 3 =>
11
               I + 1] 'Reduce ("*", 1);
12
      Put_Line ("I = "
13
                 & I'Image);
14
   end Show_Reduction_Expression;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Reduction_Expressions.

→Reduction_Expression_Value_Sequences

MD5: 25b75869e53aa3c8a8f8c821a05718c5
```

Runtime output

 $\begin{array}{rrrr} I &=& 9\\ I &=& 24 \end{array}$

In this example, we create the value sequence [for I in 1 ... 3 => I + 1] and reduce it using the + and * operators. (Note that the operations in this example have the same results as in the previous examples using arrays.)

9.5.2 Custom reducers

In the previous examples, we've used standard operators such as + and * as the reducer. We can, however, write our own reducers and pass them to the Reduce attribute. For example:

```
Listing 31: show_reduction_expression.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_Reduction_Expression is
3
       type Integer_Array is
4
         array (Positive range <>) of Integer;
5
6
      A : Integer Array (1 .. 3);
7
      I : Long_Integer;
8
9
       procedure Accumulate
10
         (Accumulator : in out Long_Integer;
11
          Value
                       : Integer) is
12
      begin
13
          Accumulator := Accumulator
14
                          + Long_Integer (Value);
15
      end Accumulate:
16
17
   begin
18
      A := [2, 3, 4];
19
       I := A'Reduce (Accumulate, 0);
20
21
```

```
22 Put_Line ("A = "
23 & & A'Image);
24 Put_Line ("I = "
25 & & I'Image);
26 end Show_Reduction_Expression;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Reduction_Expressions.

→Custom_Reducer_Procedure

MD5: 1ed7cd1f3f5d5b8acda36b04afa955f0
```

Runtime output

A = [2, 3, 4] I = 9

In this example, we implement the Accumulate procedure as our reducer, which is called to accumulate the individual elements (integer values) of the list. We pass this procedure to the Reduce attribute in the I := A'Reduce (Accumulate, 0) statement, which is equivalent to:

```
I := 0;
for E of A loop
        Accumulate (I, E);
end loop;
```

A custom reducer must have the following parameters:

- 1. The accumulator parameter, which stores the interim result and the final result as well, once all elements of the list have been processed.
- 2. The value parameter, which is a single element from the list.

Note that the accumulator type doesn't need to match the type of the individual components. In this example, we're using **Integer** as the component type, while the accumulator type is **Long_Integer**. (For this kind of reducers, using **Long_Integer** instead of **Integer** for the accumulator type makes lots of sense due to the risk of triggering overflows while the reducer is accumulating values — e.g. when accumulating a long list with larger numbers.)

In the example above, we've implemented the reducer as a procedure. However, we can also implement it as a function. In this case, the accumulated value is returned by the function:

Listing 32: show_reduction_expression.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Reduction_Expression is
3
      type Integer Array is
4
        array (Positive range <>) of Integer;
5
6
      A : Integer Array (1 .. 3);
7
      I : Long_Integer;
8
9
      function Accumulate
10
         (Accumulator : Long_Integer;
11
         Value
                     : Integer)
12
          return Long_Integer is
13
```

```
begin
14
          return Accumulator + Long_Integer (Value);
15
       end Accumulate;
16
17
   begin
18
       A := [2, 3, 4];
19
       I := A'Reduce (Accumulate, 0);
20
21
       Put_Line ("A = "
22
                  & A'Image);
23
       Put_Line ("I = "
24
                  & I'Image);
25
26
   end Show_Reduction_Expression;
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Expressions.Reduction_Expressions. →Custom_Reducer_Function MD5: 3bfc9b59e4667490e40921770990f52b

Runtime output

A = [2, 3, 4] I = 9

In this example, we converted the Accumulate procedure into a function (while the core implementation is essentially the same).

Note that the reduction expression remains the same, independently of whether we're using a procedure or a function as the reducer. Therefore, the statement with the reduction expression in this example is the same as in the previous example: I := A'Reduce (Accumulate, 0); Now that we're using a function, this statement is equivalent to:

```
I := 0;
for E of A loop
I := Accumulate (I, E);
end loop;
```

9.5.3 Other accumulator types

The accumulator type isn't restricted to scalars: in fact, we could use record types as well. For example:

```
Listing 33: show_reduction_expression.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Reduction Expression is
3
      type Integer Array is
4
        array (Positive range <>) of Integer;
5
6
      A : Integer Array (1 .. 3);
7
8
      type Integer_Accumulator is record
9
         Value : Long_Integer;
10
          Count : Integer;
11
      end record;
12
13
      function Accumulate
14
```

```
(Accumulator : Integer_Accumulator;
15
          Value
                 : Integer)
16
          return Integer_Accumulator is
17
      beain
18
          return (Value => Accumulator.Value
19
                            + Long_Integer (Value),
20
                  Count => Accumulator.Count + 1);
21
      end Accumulate;
22
23
      function Zero return Integer_Accumulator is
24
         (Value => 0, Count => 0);
25
26
      function Average (Acc : Integer_Accumulator)
27
                          return Float is
28
         (Float (Acc.Value) / Float (Acc.Count));
29
30
      Acc : Integer_Accumulator;
31
32
   begin
33
      A := [2, 3, 4];
34
35
      Acc := A'Reduce (Accumulate, Zero);
36
      Put_Line ("Acc = "
37
                 & Acc'Image);
38
      Put_Line ("Avg = "
39
                 & Average (Acc)'Image);
40
   end Show_Reduction_Expression;
41
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Reduction_Expressions.

→Reducer_Integer_Accumulator

MD5: 95d6le18e7b719d0a25dc35cdbff6af2
```

Runtime output

Acc = (VALUE => 9, COUNT => 3) Avg = 3.00000E+00

In this example, we're using the Integer_Accumulator record type in our reducer — the Accumulate function. In this case, we're not only accumulating the values, but also counting the number of elements in the list. (Of course, we could have used A'Length for that as well.)

Also, we're not limited to numeric types: we can also create a reducer using strings as the accumulator type. In fact, we can display the initial value and the elements of the list by using unbounded strings:

Listing 34: show_reduction_expression.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Strings.Unbounded;
procedure Show_Reduction_Expression is
type Integer_Array is
array (Positive range <>) of Integer;
```

```
A : Integer_Array (1 .. 3);
10
11
       function Unbounded_String_List
12
         (Accumulator : Unbounded_String;
13
          Value
                      : Integer)
14
              return Unbounded_String is
15
      begin
16
          return Accumulator
17
                 & ", " & Value'Image;
18
      end Unbounded_String_List;
19
20
   begin
21
      A := [2, 3, 4];
22
23
      Put_Line ("A = "
24
                 & A'Image);
25
      Put_Line ("L = "
26
                 \& To_String (A'Reduce
27
                    (Unbounded_String_List,
28
                       To_Unbounded_String ("0")));
29
   end Show_Reduction_Expression;
30
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Expressions.Reduction_Expressions.

⊶Reducer_String_Accumulator

MD5: 557416f08f28a48110c0fa6909086629
```

Runtime output

A = [2, 3, 4] L = 0, 2, 3, 4

In this case, the "accumulator" is concatenating the initial value and individual values of the list into a string.

STATEMENTS

10.1 Simple and Compound Statements

We can classify statements as either simple or compound. Simple statements don't contain other statements; think of them as "atomic units" that cannot be further divided. Compound statements, on the other hand, may contain other — simple or compound — statements.

Here are some examples from each category:

Category	Examples
Simple statements	Null statement, assignment, subprogram call, etc.
Compound statements	If statement, case statement, loop statement, block statement

In the Ada Reference Manual	
 5.1 Simple and Compound Statements - Sequences of Statements¹⁸² 	

10.2 Labels

We can use labels to identify statements in the code. They have the following format: <<**Some_Label>>**. We write them right before the statement we want to apply it to. Let's see an example of labels with simple statements:

Listing 1:	show	statement	identifier.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_Statement_Identifier is
3
      pragma Warnings (Off, "is not referenced");
4
   begin
5
      <<Show_Hello>> Put_Line ("Hello World!");
6
      <<Show_Test>> Put_Line ("This is a test.");
7
8
      <<Show Separator>>
9
      <<Show_Block_Separator>>
10
      Put Line ("=======");
11
  end Show_Statement_Identifier;
12
```

Code block metadata

¹⁸² http://www.ada-auth.org/standards/22rm/html/RM-5-1.html

Project: Courses.Advanced_Ada.Control_Flow.Statements.Labels.Simple_Labels
MD5: 820f5963b476af5c04314fd4373d2286

Runtime output

Hello World! This is a test.

Here, we're labeling each statement. For example, we use the Show_Hello label to identify the Put_Line ("Hello World!"); statement. Note that we can use multiple labels a single statement. In this code example, we use the Show_Separator and Show_Block_Separator labels for the same statement.

In the Ada Reference Manual

• 5.1 Simple and Compound Statements - Sequences of Statements¹⁸³

10.2.1 Labels and goto statements

Labels are mainly used in combination with **goto** statements. (Although pretty much uncommon, we could potentially use labels to indicate important statements in the code.) Let's see an example where we use a **goto** label; statement to *jump* to a specific label:

Listing 2: show c	leanup.adb
-------------------	------------

```
procedure Show Cleanup is
1
      pragma Warnings (Off, "always false");
2
3
      Some Error : Boolean;
4
   begin
5
      Some_Error := False;
6
7
      if Some Error then
8
         goto Cleanup;
9
      end if;
10
11
      <<Cleanup>> null:
12
   end Show Cleanup;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Statements.Labels.Label_Goto
MD5: 0ce06582bbefae818d4da3b7d2d3436b
```

Here, we transfer the control to the *cleanup* statement as soon as an error is detected.

10.2.2 Use-case: Continue

Another use-case is that of a Continue label in a loop. Consider a loop where we want to skip further processing depending on a condition:

Listing 3: show_continue.adb

```
procedure Show_Continue is
function Is Further Processing Needed
```

(continues on next page)

¹⁸³ http://www.ada-auth.org/standards/22rm/html/RM-5-1.html

```
(Dummy : Integer)
3
          return Boolean
4
       is
5
       begin
6
             Dummy implementation
7
          - -
          return False;
8
       end Is_Further_Processing_Needed;
9
10
       A : constant array (1 .. 10) of Integer :=
11
            (others => 0);
12
   begin
13
       for E of A loop
14
15
          -- Some stuff here...
16
17
          if Is_Further_Processing_Needed (E) then
18
19
              -- Do more stuff...
20
21
             null;
22
          end if;
23
       end loop;
24
   end Show_Continue;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Statements.Labels.Label_Continue_1
MD5: 115eeaf08d5fb072d707d6325fe9cfd0
```

In this example, we call the Is_Further_Processing_Needed (E) function to check whether further processing is needed or not. If it's needed, we continue processing in the **if** statement. We could simplify this code by just using a Continue label at the end of the loop and a **goto** statement:

Listing 4: show_continue.adb

```
procedure Show Continue is
1
       function Is_Further_Processing_Needed
2
         (Dummy : Integer)
3
          return Boolean
4
      is
5
      begin
6
          -- Dummy implementation
7
          return False;
8
      end Is Further Processing Needed;
9
10
       A : constant array (1 .. 10) of Integer :=
11
         (others => 0);
12
   begin
13
      for E of A loop
14
15
          -- Some stuff here...
16
17
          if not Is Further Processing Needed (E) then
18
             goto Continue;
19
          end if;
20
21
          -- Do more stuff...
22
23
          <<Continue>>
24
      end loop;
25
```

26 end Show_Continue;

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Statements.Labels.Label_Continue_2
MD5: 260b52ead782adf76eee5cf3c4e8332b
```

Here, we use a Continue label at the end of the loop and jump to it in the case that no further processing is needed. Note that, in this example, we don't have a statement after the Continue label because the label itself is at the end of a statement — to be more specific, at the end of the loop statement. In such cases, there's an implicit **null** statement.

```
    Historically
Since Ada 2012, we can simply write:
    loop
    -- Some statements...
    <<Continue>>
    end loop;
    If a label is used at the end of a sequence of statements, a null statement is implied.
    In previous versions of Ada, however, that is not the case. Therefore, when using those
    versions of the language, we must write at least a null statement:
    loop
        -- Some statements...
```

<<Continue>> null;

end loop;

10.2.3 Labels and compound statements

We can use labels with compound statements as well. For example, we can label a **for** loop:

```
Listing 5: show_statement_identifier.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show_Statement_Identifier is
3
      pragma Warnings (Off, "is not referenced");
4
5
             : constant array (1 .. 5) of Integer :=
6
      Arr
                 (1, 4, 6, 42, 49);
7
      Found : Boolean := False;
8
   begin
9
      <<Find 42>> for E of Arr loop
10
         if E = 42 then
11
             Found := True;
12
             exit;
13
         end if;
14
      end loop;
15
16
      Put_Line ("Found: " & Found'Image);
17
   end Show_Statement_Identifier;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Statements.Labels.Loop_Label
MD5: 5ca80b5a379ba0b08ccfaa4c6eab64d5

Runtime output

Found: TRUE

1 2

3

4

5

6

7

8

9

10

11

12

13 14

15

16

1 For further reading...

In addition to labels, loops and block statements allow us to use a statement identifier. In simple terms, instead of writing <<**Some_Label>>**, we write Some_Label :.

We could rewrite the previous code example using a loop statement identifier:

Listing 6: show statement identifier.adb

Code block metadata

Runtime output

Found: TRUE

Loop statement and block statement identifiers are generally preferred over labels. Later in this chapter, we discuss this topic in more detail.

10.3 Exit loop statement

We've introduced bare loops back in the Introduction to Ada course¹⁸⁴. In this section, we'll briefly discuss loop names and exit loop statements.

A bare loop has this form:

```
loop
    exit when Some_Condition;
end loop;
```

We can name a loop by using a loop statement identifier:

¹⁸⁴ https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html#intro-ada-bare-loops

```
Loop_Name:
    loop
    exit Loop_Name when Some_Condition;
    end loop Loop_Name;
```

In this case, we have to use the loop's name after **end loop**. Also, having a name for a loop allows us to indicate which loop we're exiting from: **exit** Loop_Name **when**.

Let's see a complete example:

Listing 7: show vector cursor iteration.adb

```
with Ada.Text_I0;
                                   use Ada.Text_I0;
1
2
   with Ada.Containers.Vectors;
3
   procedure Show_Vector_Cursor_Iteration is
4
5
      package Integer_Vectors is new
6
        Ada.Containers.Vectors
7
           (Index_Type => Positive,
8
            Element_Type => Integer);
9
10
      use Integer_Vectors;
11
12
      V : constant Vector := 20 & 10 & 0 & 13;
13
      C : Cursor;
14
15
   begin
      C := V.First;
16
      Put_Line ("Vector elements are: ");
17
18
      Show Elements :
19
         loop
20
             exit Show_Elements when C = No_Element;
21
22
             Put_Line ("Element: "
23
                        & Integer'Image (V (C)));
24
             C := Next (C);
25
         end loop Show_Elements;
26
27
   end Show_Vector_Cursor_Iteration;
28
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Statements.Exit_Loop_Statement.Exit_ →Named_Loop MD5: b77353f6ed98f8ddb32c73c47d249020

Runtime output

Vector elements are: Element: 20 Element: 10 Element: 0 Element: 13

Naming a loop is particularly useful when we have nested loops and we want to exit directly from the inner loop:

Listing 8: show_inner_loop_exit.adb

```
procedure Show_Inner_Loop_Exit is
pragma Warnings (Off);
```

```
3
      Cond : Boolean := True;
4
   begin
5
6
      Outer_Processing : loop
7
8
          Inner_Processing : loop
9
             exit Outer Processing when Cond;
10
          end loop Inner_Processing;
11
12
       end loop Outer Processing;
13
14
   end Show_Inner_Loop_Exit;
15
```

Code block metadata

Here, we indicate that we exit from the Outer_Processing loop in case a condition Cond is met, even if we're actually within the inner loop.

In the Ada Reference Manual

5.7 Exit Statements¹⁸⁵

10.4 If, case and loop statements

In the Introduction to Ada course, we talked about if statements¹⁸⁶, loop statements¹⁸⁷, and case statements¹⁸⁸. This is a very simple code example with these statements:

Listing 9: show_if_case_loop_statements.adb

```
procedure Show_If_Case_Loop_Statements is
1
      pragma Warnings (Off);
2
3
      Reset
                 : Boolean := False;
4
5
      Increment : Boolean := True;
                 : Integer := 0;
6
      Val
7
   begin
8
       -- If statement
9
10
      if Reset then
11
          Val := 0;
12
      elsif Increment then
13
          Val := Val + 1;
14
      else
15
          Val := Val - 1;
16
17
      end if;
```

```
<sup>185</sup> http://www.ada-auth.org/standards/22rm/html/RM-5-7.html
```

 ¹⁸⁶ https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html#intro-ada-if-statement
 ¹⁸⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html#

intro-ada-loop-statement

 $^{^{188}\} https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html#intro-ada-case-statement$

```
18
19
       -- Loop statement
20
21
       for I in 1 .. 5 loop
22
         Val := Val * 2 - I;
23
       end loop;
24
25
26
       - -
          Case statement
27
28
       case Val is
29
          when 0 .. 5 =>
30
31
             null;
          when others =>
32
             Val := 5;
33
       end case;
34
35
   end Show_If_Case_Loop_Statements;
36
```

Code block metadata

In this section, we'll look into a more advanced detail about the case statement.

In the Ada Reference Manual

- 5.3 If Statements¹⁸⁹
- 5.4 Case Statements¹⁹⁰
- 5.5 Loop Statements¹⁹¹

10.4.1 Case statements and expressions

As we know, the case statement has a choice expression (**case** Choice_Expression **is**), which is expected to be a discrete type. Also, this expression can be a function call or a type conversion, for example — in additional to being a variable or a constant.

As we discussed *earlier on* (page 431), if we use parentheses, the contents between those parentheses is parsed as an expression. In the context of case statements, the expression is first evaluated before being used as a choice expression. Consider the following code example:

Listing 10: scales.ads

1	puckage	States 15	
2			
3	type	Satisfaction_Scale is	(Very_Dissatisfied,
4			Dissatisfied,
5			0K,
6			Satisfied,
7			<pre>Very_Satisfied);</pre>

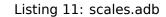
(continues on next page)

¹⁸⁹ http://www.ada-auth.org/standards/22rm/html/RM-5-3.html
¹⁹⁰ http://www.ada-auth.org/standards/22rm/html/RM-5-4.html

nackage Scales is

¹⁹¹ http://www.ada-auth.org/standards/22rm/html/RM-5-5.html

```
8
9 type Scale is range 0 .. 10;
10
11 function To_Satisfaction_Scale
12 (S : Scale)
13 return Satisfaction_Scale;
14
15 end Scales;
```



```
package body Scales is
1
2
      function To_Satisfaction_Scale
3
         (S : Scale)
4
          return Satisfaction_Scale
5
      is
6
          Satisfaction : Satisfaction_Scale;
7
      begin
8
         case (S) is
9
             when 0 .. 2 =>
10
                Satisfaction := Very_Dissatisfied;
11
             when 3 .. 4 =>
12
                Satisfaction := Dissatisfied;
13
             when 5 .. 6 =>
14
                Satisfaction := OK;
15
             when 7 ... 8 =>
16
                Satisfaction := Satisfied;
17
18
             when 9 .. 10 =>
                Satisfaction := Very_Satisfied;
19
20
         end case;
21
          return Satisfaction;
22
      end To_Satisfaction_Scale;
23
24
   end Scales;
25
```

Listing 12: show_case_statement_expression.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Scales;
                      use Scales;
3
4
5
   procedure Show_Case_Statement_Expression is
      Score : constant Scale := 0;
6
   begin
7
      Put_Line ("Score: "
8
                 & Scale'Image (Score)
9
                 & Satisfaction_Scale'Image (
10
                     To_Satisfaction_Scale (Score)));
11
12
   end Show_Case_Statement_Expression;
13
```

Code block metadata

Build output

scales.adb:9:07: error: missing case values: -128 .. -1
scales.adb:9:07: error: missing case values: 11 .. 127
gprbuild: *** compilation phase failed

When we try to compile this code example, the compiler complains about missing values in the To_Satisfaction_Scale function. As we mentioned in the Introduction to Ada course¹⁹², every possible value for the choice expression needs to be covered by a unique branch of the case statement. In principle, it *seems* that we're actually covering all possible values of the Scale type, which ranges from 0 to 10. However, we've written **case** (S) **is** instead of **case** S **is**. Because of the parentheses, (S) is evaluated as an expression. In this case, the expected range of the case statement is not Scale 'Range, but the range of its *base type* (page 377) Scale 'Base 'Range.

1 In other languages

In C, the switch-case statement requires parentheses for the choice expression:

Listing 13: main.c

```
#include <stdio.h>
```

1

2 3

4

6 7

8 9

10

11

12

13

14 15

16

```
int main(int argc, const char * argv[])
{
    int s = 0;
    switch (s)
    {
        case 0:
        case 1:
            printf("Value in the 0 -- 1 range\n");
        default:
            printf("Value > 1\n");
    }
}
```

Code block metadata

Runtime output

Value in the 0 -- 1 range Value > 1

In Ada, parentheses aren't expected in the choice expression. Therefore, we shouldn't write **case** (S) **is** in a C-like fashion — unless, of course, we really want to evaluate an expression in the case statement.

10.5 Block Statements

We've introduced block statements back in the Introduction to Ada course¹⁹³. They have this simple form:

¹⁹² https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html# intro-ada-case-statement

¹⁹³ https://learn.adacore.com/courses/intro-to-ada/chapters/imperative_language.html# intro-ada-block-statement

Listing 14: show_block_statement.adb

```
procedure Show Block Statement is
1
      pragma Warnings (Off);
2
   begin
3
4
       -- BLOCK STARTS HERE:
5
      declare
6
          I : Integer;
7
      begin
8
         I := 0;
9
      end;
10
11
   end Show_Block_Statement;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Statements.Block_Statements.Simple_

→Block_Statement

MD5: 61134b3899620c6d9ed68974fae33b5e
```

We can use an identifier when writing a block statement. (This is similar to loop statement identifiers that we discussed in the previous section.) In this example, we implement a block called Simple_Block:

Listing 15: show_block_statement.adb

```
procedure Show_Block_Statement is
1
      pragma Warnings (Off);
2
   begin
3
4
       Simple Block : declare
5
          I : Integer;
6
7
      begin
          I := 0;
8
      end Simple_Block;
9
10
   end Show_Block_Statement;
11
```

Code block metadata

Note that we must write end Simple_Block; when we use the Simple_Block identifier.

Block statement identifiers are useful:

- to indicate the begin and the end of a block as some blocks might be long or nested in other blocks;
- to indicate the purpose of the block (i.e. as code documentation).

```
    In the Ada Reference Manual
    5.6 Block Statements<sup>194</sup>
```

```
<sup>194</sup> http://www.ada-auth.org/standards/22rm/html/RM-5-6.html
```

10.6 Extended return statement

A common idiom in Ada is to build up a function result in a local object, and then return that object:

Listing 16: show_return.adb

```
procedure Show_Return is
1
2
       type Array Of Natural is
3
        array (Positive range <>) of Natural;
4
5
       function Sum (A : Array Of Natural)
6
                      return Natural
7
8
      is
          Result : Natural := 0;
9
      begin
10
          for Index in A'Range loop
11
             Result := Result + A (Index);
12
          end loop;
13
          return Result:
14
      end Sum:
15
16
   begin
17
      null;
18
  end Show Return;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Statements.Extended_Return_Statements.

Simple_Return

MD5: 16e85a8cba869802f912627c40a64c20
```

Since Ada 2005, a notation called the extended return statement is available: this allows you to declare the result object and return it as part of one statement. It looks like this:

Listing 17: show_extended_return.adb

```
procedure Show_Extended_Return is
1
2
      type Array_Of_Natural is
3
         array (Positive range <>) of Natural;
4
5
       function Sum (A : Array_Of_Natural)
6
                      return Natural
7
      is
8
      beain
9
          return Result : Natural := 0 do
10
             for Index in A'Range loop
11
                Result := Result + A (Index);
12
             end loop;
13
          end return;
14
      end Sum;
15
16
   begin
17
      null;
18
   end Show_Extended_Return;
19
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Statements.Extended_Return_Statements. ⇔Extended_Return

```
MD5: d6d6edaf800a0e346ff8ede13cbbe100
```

The return statement here creates Result, initializes it to 0, and executes the code between **do** and **end** return. When **end** return is reached, Result is automatically returned as the function result.

In the Ada Reference Manual

• 6.5 Return Statements¹⁹⁵

10.6.1 Other usages of extended return statements

\rm 1 Note

This section was originally written by Robert A. Duff and published as Gem #10: Limited Types in Ada 2005¹⁹⁶.

While the extended_return_statement was added to the language specifically to support *limited constructor functions* (page 823), it comes in handy whenever you want a local name for the function result:

Listina	18:	show	strina	_construct.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_String_Construct is
3
4
       function Make_String
5
         (S
                      : String;
6
          Prefix
                     : String;
7
8
          Use_Prefix : Boolean) return String
       is
9
          Length : Natural := S'Length;
10
      begin
11
          if Use_Prefix then
12
             Length := Length + Prefix'Length;
13
          end if;
14
15
          return Result : String (1 .. Length) do
16
17
                fill in the characters
18
             if Use_Prefix then
19
                 Result
20
                   (1 .. Prefix'Length) := Prefix;
21
22
                 Result
23
                   (Prefix'Length + 1 .. Length) := S;
24
             else
25
                 Result := S;
26
             end if;
27
28
          end return;
29
       end Make_String;
30
31
```

(continues on next page)

¹⁹⁵ http://www.ada-auth.org/standards/22rm/html/RM-6-5.html
 ¹⁹⁶ https://www.adacore.com/gems/ada-gem-10

```
S1 : String := "Ada";
32
      S2 : String := "Make_With_";
33
   begin
34
      Put_Line ("No prefix:
35
                 & Make_String (S1, S2, False));
36
      Put_Line ("With prefix:
37
                 & Make_String (S1, S2, True));
38
   end Show_String_Construct;
39
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Statements.Extended_Return_Statements. →Extended_Return_Other_Usages MD5: a2b26ceed06a0ab66aff6c2b59c02003

Runtime output

No prefix: Ada With prefix: Make_With_Ada

In this example, we first calculate the length of the string and store it in Length. We then use this information to initialize the return object of the Make_String function.

SUBPROGRAMS

11.1 Parameter Modes and Associations

In this section, we discuss some details about parameter modes and associations. First of all, as we know, parameters can be either formal or actual:

- Formal parameters are the ones we see in a subprogram declaration and implementation;
- Actual parameters are the ones we see in a subprogram call.
 - Note that actual parameters are also called *subprogram arguments* in other languages.

We define parameter associations as the connection between an actual parameter in a subprogram call and its declaration as a formal parameter in a subprogram specification or body.

1 In the Ada Reference Manual

- 6.2 Formal Parameter Modes¹⁹⁷
- 6.4.1 Parameter Associations¹⁹⁸

11.1.1 Formal Parameter Modes

We already discussed formal parameter modes in the Introduction to Ada¹⁹⁹ course:

in	Parameter can only be read, not written
out	Parameter can be written to, then read
in out	Parameter can be both read and written

As this topic was already discussed in that course — and we used parameter modes extensively in all code examples from that course —, we won't introduce the topic again here. Instead, we'll look into some of the more advanced details.

11.1.2 By-copy and by-reference

In the Introduction to Ada²⁰⁰ course, we saw that parameter modes don't correspond directly to how parameters are actually passed. In fact, an **in out** parameter could be passed by copy. For example:

¹⁹⁷ http://www.ada-auth.org/standards/22rm/html/RM-6-2.html

¹⁹⁸ http://www.ada-auth.org/standards/22rm/html/RM-6-4-1.html

¹⁹⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/subprograms.html#intro-ada-parameter-modes

²⁰⁰ https://learn.adacore.com/courses/intro-to-ada/chapters/subprograms.html#intro-ada-parameter-modes

Listing 1: check_param_passing.ads

```
with System;
procedure Check_Param_Passing
(Formal : System.Address;
Actual : System.Address);
```

Listing 2: check_param_passing.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with System.Address_Image;
2
3
   procedure Check_Param_Passing
4
     (Formal : System.Address;
5
      Actual : System.Address) is
6
7
   begin
      Put_Line ("Formal parameter at "
8
                 & System.Address_Image (Formal));
9
10
      Put_Line ("Actual parameter at "
                 & System.Address_Image (Actual));
11
      if System.Address_Image (Formal) =
12
          System.Address_Image (Actual)
13
      then
14
          Put Line
15
            ("Parameter is passed by reference.");
16
      else
17
          Put Line
18
            ("Parameter is passed by copy.");
19
      end if;
20
   end Check_Param_Passing;
21
```

Listing 3: machine_x.ads

```
with System;
1
2
   package Machine X is
3
4
      procedure Update Value
5
        (V : in out Integer;
6
         AV :
                      System.Address);
7
8
   end Machine X;
9
```

Listing 4: machine_x.adb

```
with Check Param Passing;
1
2
   package body Machine_X is
3
4
5
      procedure Update_Value
         (V : in out Integer;
6
         AV :
                      System.Address) is
7
      begin
8
         V := V + 1;
9
         Check_Param_Passing (Formal => V'Address,
10
                                Actual => AV);
11
      end Update_Value;
12
13
   end Machine_X;
14
```

Listing 5: show_by_copy_by_ref_params.adb

```
with Machine_X; use Machine_X;
procedure Show_By_Copy_By_Ref_Params is
A : Integer := 5;
begin
Update_Value (A, A'Address);
end Show_By_Copy_By_Ref_Params;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_
⇔Associations.By_Copy_By_Ref_Params
MD5: e437d3432703124496f0a217177959eb
```

Runtime output

Formal parameter at 00007FFEF3EFC6AC Actual parameter at 00007FFEF3EFC6CC Parameter is passed by copy.

As we can see by running this example,

• the integer variable A in the Show_By_Copy_By_Ref_Params procedure

and

• the V parameter in the Update_Value procedure

have different addresses, so they are different objects. Therefore, we conclude that this parameter is being passed by value, even though it has the **in out** mode. (We talk more about addresses and the 'Address attribute *later on* (page 123)).

As we know, when a parameter is passed by copy, it is first copied to a temporary object. In the case of a parameter with **in out** mode, the temporary object is copied back to the original (actual) parameter at the end of the subprogram call. In our example, the temporary object indicated by V is copied back to A at the end of the call to Update_Value.

In Ada, it's not the parameter mode that determines whether a parameter is passed by copy or by reference, but rather its type. We can distinguish between three categories:

- 1. By-copy types;
- 2. By-reference types;
- 3. Unspecified types.

Obviously, parameters of by-copy types are passed by copy and parameters of by-reference type are passed by reference. However, if a category isn't specified — i.e. when the type is neither a by-copy nor a by-reference type —, the decision is essentially left to the compiler.

As a rule of thumb, we can say that;

- elementary types and any type that is essentially elementary, such as a private type whose full view is an elementary type — are passed by copy;
- tagged and explicitly limited types and other types that are essentially tagged, such as task types — are passed by reference.

The following table provides more details:

Type category	Parameter passing	List of types
Ву сору	Ву сору	 Elementary types Descendant of a private type whose full type is a by-copy type
By reference	By reference	 Tagged types Task and protected types Explicitly limited record types Composite types with at least one subcomponent of a by-reference type Private types whose full type is a by-reference type Any descendant of the types mentioned above
Unspecified	Either by copy or by refer- ence	 Any type not men- tioned above

Note that, for parameters of limited types, only those parameters whose type is *explicitly* limited are always passed by reference. We discuss this topic in more details *in another chapter* (page 832).

Let's see an example:

Listing 6: machine_x.ads

```
with System;
1
2
   package Machine_X is
3
4
      type Integer_Array is
5
        array (Positive range <>) of Integer;
6
7
      type Rec is record
8
         A : Integer;
9
      end record;
10
11
      type Rec_Array is record
12
         A : Integer;
13
         Arr : Integer_Array (1 .. 100);
14
      end record;
15
16
      type Tagged_Rec is tagged record
17
         A : Integer;
18
      end record;
19
20
      procedure Update_Value
21
         (R : in out Rec;
22
```

```
AR :
                       System.Address);
23
24
       procedure Update_Value
25
         (RA : in out Rec_Array;
26
          ARA :
                        System.Address);
27
28
       procedure Update_Value
29
         (R : in out Tagged_Rec;
30
          AR :
                       System.Address);
31
32
   end Machine_X;
33
```



```
with Check_Param_Passing;
1
2
   package body Machine X is
3
4
       procedure Update_Value
5
         (R : in out Rec;
6
                       System.Address
          AR :
7
       is
8
       begin
9
          R.A := R.A + 1;
10
          Check_Param_Passing (Formal => R'Address,
11
                                  Actual \Rightarrow AR);
12
       end Update_Value;
13
14
       procedure Update_Value
15
         (RA : in out Rec_Array;
16
                        System.Address
17
          ARA :
       is
18
       begin
19
          RA.A := RA.A + 1;
20
          Check_Param_Passing (Formal => RA'Address,
21
                                  Actual => ARA);
22
       end Update_Value;
23
24
       procedure Update Value
25
         (R : in out Tagged_Rec;
26
27
          AR :
                       System.Address)
       is
28
       begin
29
          R.A := R.A + 1;
30
          Check_Param_Passing (Formal => R'Address,
31
                                  Actual \Rightarrow AR);
32
       end Update_Value;
33
34
   end Machine X;
```

Listing 8: show by copy by ref params.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Machine_X;
                     use Machine_X;
3
4
   procedure Show_By_Copy_By_Ref_Params is
      TR : Tagged_Rec := (A
                             => 5);
5
                              => 5);
      R : Rec
                      := (A
6
                      := (A
                               => 5,
      RA : Rec_Array
7
                           Arr => (others => 0));
8
```

35

```
begin
9
      Put_Line ("Tagged record");
10
      Update_Value (TR, TR'Address);
11
12
      Put Line ("Untagged record");
13
      Update_Value (R, R'Address);
14
15
      Put Line ("Untagged record with array");
16
      Update Value (RA, RA'Address);
17
   end Show_By_Copy_By_Ref_Params;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_
⇔Associations.By_Copy_By_Ref_Params
MD5: 3ca46380c4df36af9393041181ff2f17
```

Runtime output

```
Tagged record
Formal parameter at 00007FFC27865450
Actual parameter at 00007FFC27865450
Parameter is passed by reference.
Untagged record
Formal parameter at 00007FFC2786529C
Actual parameter at 00007FFC2786544C
Parameter is passed by copy.
Untagged record with array
Formal parameter at 00007FFC278652B0
Actual parameter at 00007FFC278652B0
Parameter is passed by reference.
```

When we run this example, we see that the object of tagged type (Tagged_Rec) is passed by reference to the Update_Value procedure. In the case of the objects of untagged record types, you might see this:

- the parameter of Rec type which is an untagged record with a single component of integer type —, the parameter is passed by copy;
- the parameter of Rec_Array type which is an untagged record with a large array of 100 components —, the parameter is passed by reference.

Because Rec and Rec_Array are neither by-copy nor by-reference types, the decision about how to pass them to the Update_Value procedure is made by the compiler. (Thus, it is possible that you see different results when running the code above.)

11.1.3 Bounded errors

When we use parameters of types that are neither by-copy nor by-reference types, we might encounter the situation where we have the same object bound to different names in a subprogram. For example, if:

we use a global object Global_R of a record type Rec

and

• we have a subprogram with an in-out parameter of the same record type Rec

and

 we pass Global_R as the actual parameter for the in-out parameter of this subprogram, then we have two access paths to this object: one of them using the global variable directly, and the other one using it indirectly via the in-out parameter. This situation could lead to undefined behavior or to a program error. Consider the following code example:

```
Listing 9: machine_x.ads
```

```
with System;
1
2
   package Machine X is
3
4
       type Rec is record
5
          A : Integer;
6
       end record;
7
8
       Global_R : Rec := (A \Rightarrow 0);
9
10
       procedure Update_Value
11
         (R : in out Rec;
12
          AR :
                      System.Address);
13
14
   end Machine_X;
15
```

Listing 10: machine_x.adb

```
use Ada.Text I0;
   with Ada.Text I0;
1
2
   with Check Param Passing;
3
4
   package body Machine X is
5
6
      procedure Update Value
7
8
         (R : in out Rec;
          AR :
                    System.Address)
9
      is
10
          procedure Show Vars is
11
          beain
12
             Put Line ("Global R.A: "
13
                        & Integer'Image (Global_R.A));
14
             Put Line ("R.A:
15
                        & Integer'Image (R.A));
16
          end Show_Vars;
17
      begin
18
          Check_Param_Passing (Formal => R'Address,
19
                                 Actual => AR);
20
21
          Put Line ("Incrementing Global R.A...");
22
          Global R.A := Global R.A + 1;
23
          Show Vars;
24
25
          Put Line ("Incrementing R.A...");
26
          R.A := R.A + 5;
27
          Show Vars;
28
      end Update Value;
29
30
```

31 end Machine_X;

Listing 11: show_by_copy_by_ref_params.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Machine_X; use Machine_X;
procedure Show_By_Copy_By_Ref_Params is
```

```
5 begin
6 Put_Line ("Calling Update_Value...");
7 Update_Value (Global_R, Global_R'Address);
8
9 Put_Line ("After call to Update_Value...");
10 Put_Line ("Global_R.A: "
11 & & Integer'Image (Global_R.A));
12 end Show_By_Copy_By_Ref_Params;
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_ ⇔Associations.By_Copy_By_Ref_Params MD5: 96be7054b7ff64a304705edf6b15f031

Runtime output

```
Calling Update_Value...
Formal parameter at 00007FFD7131E05C
Actual parameter at 0000591CC4E103C4
Parameter is passed by copy.
Incrementing Global_R.A...
Global_R.A: 1
R.A: 0
Incrementing R.A...
Global_R.A: 1
R.A: 5
After call to Update_Value...
Global_R.A: 5
```

In the Update_Value procedure, because Global_R and R have a type that is neither a by-pass nor a by-reference type, the language does not specify whether the old or the new value would be read in the calls to Put_Line. In other words, the actual behavior is undefined. Also, this situation might raise the Program_Error exception.

Important

As a general advice:

- you should be very careful when using global variables and
- you should avoid passing them as parameters in situations such as the one illustrated in the code example above.

11.1.4 Aliased parameters

When a parameter is specified as *aliased*, it is always passed by reference, independently of the type we're using. In this sense, we can use this keyword to circumvent the rules mentioned so far. (We discuss more about *aliasing* (page 634) and *aliased parameters* (page 643) later on.)

Let's rewrite a previous code example that has a parameter of elementary type and change it to *aliased*:

Listing 12: machine_x.ads

```
with System;
```

```
package Machine X is
```

(continues on next page)

2

```
4
5 procedure Update_Value
6 (V : aliased in out Integer;
7 AV : System.Address);
8
9 end Machine_X;
```

Listing 13: machine_x.adb

```
with Check_Param_Passing;
1
2
   package body Machine X is
3
4
      procedure Update Value
5
         (V : aliased in out Integer;
6
          AV :
                               System.Address)
7
      is
8
      begin
9
          V := V + 1;
10
          Check Param Passing (Formal => V'Address,
11
                                Actual => AV);
12
      end Update_Value;
13
14
   end Machine_X;
15
```

Listing 14: show_by_copy_by_ref_params.adb

```
with Machine_X; use Machine_X;
procedure Show_By_Copy_By_Ref_Params is
A : aliased Integer := 5;
begin
Update_Value (A, A'Address);
end Show_By_Copy_By_Ref_Params;
```

Code block metadata

Runtime output

Formal parameter at 00007FFE2859312C Actual parameter at 00007FFE2859312C Parameter is passed by reference.

As we can see, A is now passed by reference.

Note that we can only pass aliased objects to aliased parameters. If we try to pass a nonaliased object, we get a compilation error:

Listing 15: show_by_copy_by_ref_params.adb

```
with Machine_X; use Machine_X;
procedure Show_By_Copy_By_Ref_Params is
A : Integer := 5;
begin
Update_Value (A, A'Address);
end Show_By_Copy_By_Ref_Params;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_

⇔Associations.By_Copy_By_Ref_Params

MD5: 9e6586e0b771de68040131cae81799b8
```

Build output

```
show_by_copy_by_ref_params.adb:6:18: error: actual for aliased formal "V" must be_

aliased object
gprbuild: *** compilation phase failed
```

Again, we discuss more about *aliased parameters* (page 643) and *aliased objects* (page 636) later on in the context of access types.

11.1.5 Parameter Associations

When actual parameters are associated with formal parameters, some rules are checked. As a typical example, the type of each actual parameter must match the type of the corresponding actual parameter. In this section, we see some details about how this association is made and some of the potential errors.

```
1 In the Ada Reference Manual
```

• 6.4.1 Parameter Associations²⁰¹

Parameter order and association

As we already know, when calling subprograms, we can use positional or named parameter association — or a mixture of both. Also, parameters can have default values. Let's see some examples:

Listing 16: operations.ads

```
package Operations is
procedure Add (Left : in out Integer;
Right : Float := 1.0);
end Operations;
```

Listing 17: operations.adb

```
package body Operations is
procedure Add (Left : in out Integer;
Right : Float := 1.0) is
begin
Left := Left + Integer (Right);
end Add;
end Operations;
```

Listing 18: show_param_association.adb

```
with Operations; use Operations;
2
```

(continues on next page)

²⁰¹ http://www.ada-auth.org/standards/22rm/html/RM-6-4-1.html

```
procedure Show_Param_Association is
3
      A : Integer := 5;
4
   begin
5
           Positional association
6
      Add (A, 2.0);
7
8
       -- Positional association
9
           (using default value)
       - -
10
      Add (A);
11
12
       -- Named association
13
      Add (Left => A,
14
            Right => 2.0);
15
16
       -- Named association (inversed order)
17
      Add (Right \Rightarrow 2.0,
18
            Left => A);
19
20
       -- Mixed positional / named association
21
       Add (A, Right => 2.0);
22
   end Show_Param_Association;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_

⇔Associations.Param_Association_1

MD5: 64d3f44ac2bf72317fae22658f6d218e
```

This code snippet has examples of positional and name parameter association. Also, it has an example of mixed positional / named parameter association. In most cases, the actual A parameter is associated with the formal Left parameter, and the actual 2.0 parameter is associated with the formal Right parameter.

In addition to that, parameters can have default values, so, when we write Add (A), the A variable is associated with the Left parameter and the default value (1.0) is associated with the Right parameter.

Also, when we use named parameter association, the parameter order is irrelevant: we can, for example, write the last parameter as the first one. Therefore, we can write Add (Right => 2.0, Left => A) instead of Add (Left => A, Right => 2.0).

Ambiguous calls

Ambiguous calls can be detected by the compiler during parameter association. For example, when we have both default values in parameters and subprogram overloading, the compiler might be unable to decide which subprogram we're calling:

Listing 19: o	perations.a	ds
---------------	-------------	----

```
package Operations is
procedure Add (Left : in out Integer);
procedure Add (Left : in out Integer;
Right : Float := 1.0);
end Operations;
```

Listing 20: operations.adb

```
package body Operations is
1
2
      procedure Add (Left : in out Integer) is
3
      begin
4
         Left := Left + 1;
5
      end Add;
6
7
      procedure Add (Left : in out Integer;
8
                      Right :
                                     Float := 1.0) is
9
      begin
10
         Left := Left + Integer (Right);
11
      end Add:
12
13
   end Operations;
14
```

Listing 21: show_param_association.adb

```
with Operations; use Operations;
1
2
  procedure Show Param Association is
3
      A : Integer := 5;
4
  begin
5
      Add (A);
6
      -- ERROR: cannot decide which
7
                 procedure to take
      - -
8
  end Show_Param_Association;
9
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_

⇔Associations.Param_Association_1

MD5: 2725517f82d4068b669028eca1815079
```

Build output

```
show_param_association.adb:6:04: error: ambiguous expression (cannot resolve "Add")
show_param_association.adb:6:04: error: possible interpretation at operations.ads:5
show_param_association.adb:6:04: error: possible interpretation at operations.ads:3
gprbuild: *** compilation phase failed
```

As we see in this example, the Add procedure is overloaded. The first instance has one parameter, and the second instance has two parameters, where the second parameter has a default value. When we call Add with just one parameter, the compiler cannot decide whether we intend to call

the first instance of Add with one parameter

or

• the second instance of Add using the default value for the second parameter.

In this specific case, there are multiple options to solve the issue, but all of them involve redesigning the package specification:

- we could just rename one of Add procedures (thereby eliminating the subprogram overloading);
- we could rename the first parameter of one of the Add procedures and use named parameter association in the call to the procedure;
 - For example, we could rename the parameter to Value and call Add (Value => A).

• remove the default value from the second parameter of the second instance of Add.

Overlapping actual parameters

When we have more than one **out** or **in out** parameters in a subprogram, we might run into the situation where the actual parameter overlaps with another parameter. For example:

Listing 22: machine x.ads

```
package Machine_X is
procedure Update_Value (V1 : in out Integer;
V2 : out Integer);
end Machine_X;
```

```
Listing 23: machine_x.adb
```

```
package body Machine_X is
1
2
      procedure Update_Value (V1 : in out Integer;
3
                                V2 :
                                     out Integer) is
4
      begin
5
         V1 := V1 + 1;
6
         V2 := V2 + 1;
7
      end Update Value;
8
9
   end Machine_X;
10
```

Listing 24: show by copy by ref params.adb

```
with Machine_X; use Machine_X;

procedure Show_By_Copy_By_Ref_Params is
A : Integer := 5;
begin
Update_Value (A, A);
end Show_By_Copy_By_Ref_Params;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Parameter_Modes_

⇔Associations.Illegal_Calls

MD5: d18a7056463fee9298dd1fdef0a31daf
```

Build output

In this case, we're using A for both output parameters in the call to Update_Value. Passing one variable to more than one output parameter in a given call is forbidden in Ada, so this triggers a compilation error. Depending on the specific context, you could solve this issue by using temporary variables for the other output parameters.

11.2 Operators

Operators are commonly used for variables of scalar types such as **Integer** and **Float**. In these cases, they replace *usual* function calls. (To be more precise, operators are function calls, but written in a different format.) For example, we simply write A := A + B + C; when we want to add three integer variables. A hypothetical, non-intuitive version of this operation could be A := Add (Add (A, B), C);. In such cases, operators allow for expressing function calls in a more intuitive way.

Many primitive operators exist for scalar types. We classify them as follows:

Category	Operators
Logical	and, or, xor
Relational	=, /=, <, <=, >, >=
Unary adding	+, -
Binary adding	+, -, &
Multiplying	*,/, mod , rem
Highest precedence	**, abs, not

1 In the Ada Reference Manual

• 4.5 Operators and Expression Evaluation²⁰²

11.2.1 User-defined operators

For non-scalar types, not all operators are defined. For example, it wouldn't make sense to expect a compiler to include an addition operator for a record type with multiple components. Exceptions to this rule are the equality and inequality operators (= and /=), which are defined for any type (be it scalar, record types, and array types).

For array types, the concatenation operator (δ) is a primitive operator:

Listing 25: integer_arrays.ads

```
1 package Integer_Arrays is
2
3 type Integer_Array is
4 array (Positive range <>) of Integer;
5
6 end Integer_Arrays;
```

Listing 26: show_array_concatenation.adb

```
use Ada.Text I0;
   with Ada.Text IO;
1
   with Integer_Arrays; use Integer_Arrays;
2
3
   procedure Show Array Concatenation is
4
      A, B : Integer_Array (1 .. 5);
5
      R
            : Integer_Array (1 .. 10);
6
   beain
7
      A := (1 \& 2 \& 3 \& 4 \& 5);
8
      B := (6 \& 7 \& 8 \& 9 \& 10);
9
      R := A \& B;
10
11
```

(continues on next page)

²⁰² http://www.ada-auth.org/standards/22rm/html/RM-4-5.html

```
12 for E of R loop
13 Put (E'Image & ' ');
14 end loop;
15 New_Line;
16 end Show_Array_Concatenation;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Operators.Integer_Arrays_

GConcat

MD5: 1899e66ec1d0b36b10d8b89fc2dfac0e
```

Runtime output

1 2 3 4 5 6 7 8 9 10

In this example, we're using the primitive & operator to concatenate the A and B arrays in the assignment to R. Similarly, we're concatenating individual components (integer values) to create an aggregate that we assign to A and B.

In contrast to this, the addition operator is not available for arrays:

Listing 27: integer_arrays.ads

```
1 package Integer_Arrays is
2
3 type Integer_Array is
4 array (Positive range <>) of Integer;
5
6 end Integer_Arrays;
```

Listing 28: show array addition.adb

```
use Ada.Text IO:
   with Ada.Text IO:
1
   with Integer_Arrays; use Integer_Arrays;
2
3
   procedure Show Array Addition is
4
      A, B, R : Integer Array (1 .. 5);
5
   begin
6
      A := (1 \& 2 \& 3 \& 4 \& 5);
7
      B := (6 \& 7 \& 8 \& 9 \& 10);
8
      R := A + B;
9
10
      for E of R loop
11
          Put (E'Image & ' ');
12
      end loop;
13
      New Line;
14
15
   end Show_Array_Addition;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Operators.Integer_Arrays_

Addition

MD5: d94f9791523359d390a7cafd900d1268
```

Build output

We can, however, define *custom* operators for any type. For example, if a specific type doesn't have a predefined addition operator, we can define our own + operator for it.

Note that we're limited to the operator symbols that are already defined by the Ada language (see the previous table for the complete list of operators). In other words, the operator we define must be selected from one of those existing symbols; we cannot use new symbols for custom operators.

1 In other languages

Some programming languages — such as Haskell — allow you to define and use custom operator symbols. For example, in Haskell, you can create a new "broken bar" (¦) operator for integer values:

```
(;) :: Int -> Int -> Int
a ; b = a + a + b
main = putStrLn $ show (2 ; 3)
This is not possible in Ada.
```

Let's define a custom addition operator that adds individual components of the Integer_Array type:

Listing 29: integer_arrays.ads

```
package Integer_Arrays is
1
2
      type Integer_Array is
3
        array (Positive range <>) of Integer;
4
5
      function "+" (Left, Right : Integer_Array)
6
                     return Integer_Array
7
        with Post =>
8
           (for all I in "+"'Result'Range =>
9
              "+"'Result (I) = Left (I) + Right (I));
10
11
   end Integer Arrays;
12
```

Listing 30: integer_arrays.adb

```
package body Integer_Arrays is
1
2
      function "+" (Left, Right : Integer_Array)
3
                     return Integer_Array
4
5
      is
         R : Integer Array (Left'Range);
6
      begin
7
          for I in Left'Range loop
8
             R(I) := Left(I) + Right(I);
9
         end loop;
10
11
         return R;
12
      end "+";
13
14
   end Integer_Arrays;
15
```

Listing 31: show_array_addition.adb

with Ada.Text_I0; use Ada.Text_I0;

```
with Integer_Arrays; use Integer_Arrays;
2
3
   procedure Show_Array_Addition is
4
       A, B, R : Integer Array (1 .. 5);
5
   begin
6
       A := (1 \& 2 \& 3 \& 4 \& 5);
7
       B := (6 \& 7 \& 8 \& 9 \& 10);
8
       R := A + B;
9
10
       for E of R loop
11
          Put (E'Image & ' ');
12
       end loop;
13
14
       New_Line;
15
   end Show_Array_Addition;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Operators.Integer_Arrays_

⇔Addition

MD5: 6f50fa47270d97d3fb50379b6275777d
```

Runtime output

7 9 11 13 15

Now, the R := A + B line doesn't trigger a compilation error anymore because the + operator is defined for the Integer_Array type.

In the implementation of the +, we return an array with the range of the Left array where each component is the sum of the Left and Right arrays. In the declaration of the + operator, we're defining the expected behavior in the postcondition. Here, we're saying that, for each index of the resulting array (**for all I in** "+" 'Result 'Range), the value of each component of the resulting array at that specific index is the sum of the components from the Left and Right arrays at the same index ("+" 'Result (I) = Left (I) + Right (I)). (**for all** denotes a *quantified expression* (page 436).)

Note that, in this implementation, we assume that the range of Right is a subset of the range of Left. If that is not the case, the Constraint_Error exception will be raised at runtime in the loop. (You can test this by declaring B as Integer_Array $(5 \ldots 10)$, for example.)

We can also define custom operators for record types. For example, we could declare two + operators for a record containing the name and address of a person:

Listing 32:	addresses.ads
-------------	---------------

```
package Addresses is
1
2
      type Person is private;
3
4
      function "+" (Name
                            : String;
5
                     Address : String)
6
                      return Person;
7
      function "+" (Left, Right : Person)
8
                      return Person;
9
10
      procedure Display (P : Person);
11
12
   private
13
14
```

```
subtype Name_String is String (1 .. 40);
15
       subtype Address_String is String (1 .. 100);
16
17
      type Person is record
18
               : Name_String;
         Name
19
         Address : Address_String;
20
      end record;
21
22
   end Addresses;
23
```

```
with Ada.Strings.Fixed; use Ada.Strings.Fixed;
1
   with Ada.Text_I0;
                             use Ada.Text_I0;
2
3
   package body Addresses is
4
5
       function "+" (Name
                            : String;
6
                      Address : String)
7
                      return Person
8
       is
9
       begin
10
          return (Name
11
                           =>
                     Head (Name,
12
                           Name String'Length),
13
                   Address =>
14
                     Head (Address,
15
                           Address_String'Length));
16
       end "+";
17
18
       function "+" (Left, Right : Person)
19
                      return Person
20
       is
21
       beain
22
          return (Name
                           => Left.Name,
23
                  Address => Right.Address);
24
       end "+";
25
26
       procedure Display (P : Person) is
27
       begin
28
                               " & P.Name);
          Put Line ("Name:
29
          Put_Line ("Address: " & P.Address);
30
          New_Line;
31
       end Display;
32
33
   end Addresses;
34
```

Listing 34: show_address_addition.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Addresses; use Addresses;
2
3
   procedure Show_Address_Addition is
4
5
      John : Person := "John" + "4 Main Street";
      Jane : Person := "Jane" + "7 High Street";
6
7
   begin
      Display (John);
8
      Display (Jane);
9
      Put_Line ("-----
                          ----");
10
11
```

Jane := Jane + John; 12 Display (Jane); 13 14

end Show_Address_Addition;

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Operators.Rec_Operator MD5: c69ff43ed5a80a0c62bad87eada14301

Runtime output

Name: John Address: 4 Main Street Name: Jane Address: 7 High Street Name: Jane Address: 4 Main Street

In this example, the first + operator takes two strings — with the name and address of a person — and returns an object of Person type. We use this operator to initialize the John and Jane variables.

The second + operator in this example brings two people together. Here, the person on the left side of the + operator moves to the home of the person on the right side. In this specific case, Jane is moving to John's house.

As a small remark, we usually expect that the + operator is commutative. In other words, changing the order of the elements in the operation doesn't change the result. However, in our definition above, this is not the case, as we can confirm by comparing the operation in both orders:

Listing 35: show address addition.adb

```
with Ada.Text IO; use Ada.Text IO;
1
                    use Addresses;
   with Addresses;
2
3
   procedure Show Address Addition is
4
      John : constant Person :=
5
                "John" + "4 Main Street";
6
      Jane : constant Person :=
7
                "Jane" + "7 High Street";
8
   begin
9
      if Jane + John = John + Jane then
10
         Put Line ("It's commutative!");
11
      else
12
         Put Line ("It's not commutative!");
13
      end if;
14
  end Show_Address_Addition;
15
```

Code block metadata

Project: Courses.Advanced Ada.Control Flow.Subprograms.Operators.Rec Operator MD5: 2af6e1a31100a1d0fa786d42cc93c09b

Runtime output

It's not commutative!

In this example, we're using the primitive = operator for the Person to assess whether the result of the addition is commutative.

1 In the Ada Reference Manual

6.1 Subprogram Declarations²⁰³

11.3 Expression functions

Usually, we implement Ada functions with a construct like this: **begin return** X; **end**;. In other words, we create a **begin** ... **end**; block and we have at least one **return** statement in that block. An expression function, in contrast, is a function that is implemented with a simple expression in parentheses, such as (X);. In this case, we don't use a **begin** ... **end**; block or a **return** statement.

As an example of an expression, let's say we want to implement a function named Is_Zero that checks if the value of the integer parameter I is zero. We can implement this function with the expression I = 0. In the usual approach, we would create the implementation by writing **is begin return** I = 0; **end** Is_Zero;. When using expression functions, however, we can simplify the implementation by just writing **is** (I = 0);. This is the complete code of Is_Zero using an expression function:

Listing 36: expr_func.ads

```
1 package Expr_Func is
2
3 function Is_Zero (I : Integer)
4 return Boolean is
5 (I = 0);
6
7 end Expr_Func;
```

Code block metadata

An expression function has the same effect as the usual version using a block. In fact, the code above is similar to this implementation of the Is_Zero function using a block:

```
Listing 37: expr_func.ads
```

```
1 package Expr_Func is
2
3 function Is_Zero (I : Integer)
4 return Boolean;
5
6 end Expr_Func;
```

Listing 38: expr func.adb

```
1 package body Expr_Func is
2
3 function Is_Zero (I : Integer)
4 return Boolean is
```

(continues on next page)

²⁰³ http://www.ada-auth.org/standards/22rm/html/RM-6-1.html

```
5 begin
6 return I = 0;
7 end Is_Zero;
8
9 end Expr_Func;
```

Code block metadata

The only difference between these two versions of the Expr_Func packages is that, in the first version, the package specification contains the implementation of the Is_Zero function, while, in the second version, the implementation is in the body of the Expr_Func package.

An expression function can be, at same time, the specification and the implementation of a function. Therefore, in the first version of the Expr_Func package above, we don't have a separate implementation of the Is_Zero function because (I = 0) is the actual implementation of the function. Note that this is only possible for expression functions; you cannot have a function implemented with a block in a package specification. For example, the following code is wrong and won't compile:

Listing	39:	expr	func.ads
---------	-----	------	----------

```
package Expr_Func is
1
2
      function Is Zero (I : Integer)
3
                          return Boolean is
4
      begin
5
6
         return I = 0;
7
      end Is_Zero;
8
  end Expr_Func;
9
```

Code block metadata

We can, of course, separate the function declaration from its implementation as an expression function. For example, we can rewrite the first version of the Expr_Func package and move the expression function to the body of the package:

```
Listing 40: expr_func.ads
```

```
package Expr_Func is
function Is_Zero (I : Integer)
feed Expr_Func;
```

Listing 41: expr func.adb

```
package body Expr_Func is
function Is Zero (I : Integer)
```

```
4 return Boolean is
5 (I = 0);
6
7 end Expr_Func;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Expression_Functions.Simple_

→Expression_Function_4

MD5: 491a491da92636a35579f870969aaf08
```

In addition, we can use expression functions in the private part of a package specification. For example, the following code declares the Is_Valid function in the specification of the My_Data package, while its implementation is an expression function in the private part of the package specification:

Listing 42: my_data.ads

```
package My Data is
1
2
       type Data is private;
3
4
       function Is_Valid (D : Data)
5
                            return Boolean;
6
7
   private
8
9
       type Data is record
10
          Valid : Boolean;
11
       end record;
12
13
       function Is_Valid (D : Data)
14
                            return Boolean is
15
         (D.Valid);
16
17
   end My_Data;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Expression_Functions.

→Private_Expression_Function_1

MD5: beb57eca67b3954097e0f7ac00ea70c9
```

Naturally, we could write the function implementation in the package body instead:

```
Listing 43: my_data.ads
```

```
package My_Data is
1
2
       type Data is private;
3
4
       function Is Valid (D : Data)
5
                            return Boolean;
6
7
   private
8
9
       type Data is record
10
          Valid : Boolean:
11
       end record;
12
13
   end My_Data;
14
```

Listing 44: my_data.adb

```
1 package body My_Data is
2
3 function Is_Valid (D : Data)
4 return Boolean is
5 (D.Valid);
6
7 end My_Data;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Expression_Functions.

⊶Private_Expression_Function_2

MD5: 3c6e2a3c53c7c8e1a7b86efccdc3bf8d
```

1 In the Ada Reference Manual

• 6.8 Expression functions²⁰⁴

11.4 Overloading

\rm 1 Note

This section was originally written by Robert A. Duff and published as Gem #50: Overload Resolution²⁰⁵.

Ada allows overloading of subprograms, which means that two or more subprogram declarations with the same name can be visible at the same place. Here, "name" can refer to operator symbols, like "+". Ada also allows overloading of various other notations, such as literals and aggregates.

In most languages that support overloading, overload resolution is done "bottom up" — that is, information flows from inner constructs to outer constructs. As usual, computer folks draw their trees upside-down, with the root at the top. For example, if we have two procedures Print:

Listing 45: show overloading.adb

```
procedure Show_Overloading is
1
2
      package Types is
3
          type Sequence is null record;
4
          type Set is null record;
5
6
          procedure Print (S : Sequence) is null;
7
          procedure Print (S : Set) is null;
8
      end Types;
9
10
      use Types;
11
12
      X : Sequence;
13
   begin
14
```

(continues on next page)

²⁰⁴ http://www.ada-auth.org/standards/22rm/html/RM-6-8.html
 ²⁰⁵ https://www.adacore.com/gems/gem-50

```
15
16 -- Compiler selects Print (S : Sequence)
17 Print (X);
18 end Show_Overloading;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Overloading.Overloading
MD5: 020c4f04285c80c1050d8edbaf2dbcae
```

the type of X determines which Print is meant in the call.

Ada is unusual in that it supports top-down overload resolution as well:

Listing 46: show_top_down_overloading.adb

```
procedure Show_Top_Down_Overloading is
1
2
      package Types is
3
          type Sequence is null record;
4
          type Set is null record;
5
6
          function Empty return Sequence is
7
            ((others => <>));
8
9
          function Empty return Set is
10
            ((others => <>));
11
12
          procedure Print_Sequence (S : Sequence) is
13
            null;
14
15
          procedure Print Set (S : Set) is
16
            null;
17
18
      end Types;
19
      use Types;
20
21
      X : Sequence;
22
   begin
23
          Compiler selects function
24
       - -
          Empty return Sequence
25
      Print_Sequence (Empty);
26
   end Show_Top_Down_Overloading;
27
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Overloading.Overloading
MD5: 3b776a3efdee3d7e583ddbf5159c9a1b
```

The type of the formal parameter S of Print_Sequence determines which Empty is meant in the call. In C++, for example, the equivalent of the Print (X) example would resolve, but the Print_Sequence (Empty) would be illegal, because C++ does not use top-down information.

If we overload things too heavily, we can cause ambiguities:

Listing 47: show_overloading_error.adb

```
procedure Show_Overloading_Error is
package Types is
```

```
type Sequence is null record;
4
          type Set is null record;
5
6
          function Empty return Sequence is
7
            ((others => <>));
8
9
          function Empty return Set is
10
            ((others => <>));
11
12
          procedure Print (S : Sequence) is
13
            null;
14
15
          procedure Print (S : Set) is
16
17
            null:
       end Types;
18
19
      use Types;
20
21
      X : Sequence;
22
   begin
23
       Print (Empty); -- Illegal!
24
   end Show_Overloading_Error;
25
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Overloading.Overloading MD5: 5182c517a1afff4568ab2404ac66fda8

Build output

The call is ambiguous, and therefore illegal, because there are two possible meanings. One way to resolve the ambiguity is to use a qualified expression to say which type we mean:

Print (Sequence'(Empty));

Note that we're now using both bottom-up and top-down overload resolution: Sequence' determines which Empty is meant (top down) and which Print is meant (bottom up). You can qualify an expression, even if it is not ambiguous according to Ada rules — you might want to clarify the type because it might be ambiguous for human readers.

Of course, you could instead resolve the Print (Empty) example by modifying the source code so the names are unique, as in the earlier examples. That might well be the best solution, assuming you can modify the relevant sources. Too much overloading can be confusing. How much is "too much" is in part a matter of taste.

Ada really needs to have top-down overload resolution, in order to resolve literals. In some languages, you can tell the type of a literal by looking at it, for example appending L (letter el) means "the type of this literal is long int". That sort of kludge won't work in Ada, because we have an open-ended set of integer types:

Listing 48: show_literal_resolution.adb

```
procedure Show_Literal_Resolution is

type Apple_Count is range 0 .. 100;

procedure Peel (Count : Apple_Count) is null;
begin
Peel (20);
end Show_Literal_Resolution;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Overloading.Literal_

⊲Resolution

MD5: f428b6b4c642c44ede6bc21e7522c532
```

You can't tell by looking at the literal 20 what its type is. The type of formal parameter **Count** tells us that 20 is an Apple_Count, as opposed to some other type, such as Standard. **Long_Integer**.

Technically, the type of 20 is universal_integer, which is implicitly converted to Apple_Count — it's really the result type of that implicit conversion that is at issue. But that's an obscure point — you won't go *too* far wrong if you think of the integer literal notation as being overloaded on all integer types.

Developers sometimes wonder why the compiler can't resolve something that seems obvious. For example:

Listing 49: show_literal_resolution_error.adb

```
procedure Show_Literal_Resolution_Error is
1
2
      type Apple Count is range 0 .. 100;
3
      procedure Slice (Count : Apple_Count) is null;
4
5
      type Orange Count is range 0 .. 10 000;
6
      procedure Slice (Count : Orange Count) is null;
7
   begin
8
      Slice (Count => (10 000)); -- Illegal!
9
   end Show Literal Resolution Error;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Overloading.Literal_

⊶Resolution_Error

MD5: 4789d8eea9b82649ba8e453bb861688a
```

Build output

This call is ambiguous, and therefore illegal. But why? Clearly the developer must have meant the <code>Orange_Count</code> one, because $10_{-}000$ is out of range for <code>Apple_Count</code>. And all the relevant expressions happen to be static.

Well, a good rule of thumb in language design (for languages with overloading) is that the overload resolution rules should not be "too smart". We want this example to be illegal to

avoid confusion on the part of developers reading the code. As usual, a qualified expression fixes it:

Slice (Count => Orange_Count'(10_000));

Another example, similar to the literal, is the aggregate. Ada uses a simple rule: the type of an aggregate is determined top down (i.e., from the context in which the aggregate appears). Bottom-up information is not used; that is, the compiler does not look inside the aggregate in order to determine its type.

Listing 50: show record resolution error.adb

```
procedure Show_Record_Resolution_Error is
1
2
      type Complex is record
3
         Re, Im : Float:
4
      end record;
5
6
      procedure Grind (X : Complex) is null;
7
      procedure Grind (X : String) is null;
8
   begin
9
      Grind (X => (Re => 1.0, Im => 1.0));
10
11
       - -
          Illegal!
12
   end Show_Record_Resolution_Error;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Overloading.Record_

←Resolution_Error

MD5: e3dd1f1d0c403bcf672f4bab881b8ef9
```

Build output

There are two Grind procedures visible, so the type of the aggregate could be Complex or **String**, so it is ambiguous and therefore illegal. The compiler is not required to notice that there is only one type with components Re and Im, of some real type — in fact, the compiler is not *allowed* to notice that, for overloading purposes.

We can qualify as usual:

Grind (X => Complex'(Re => 1.0, Im => 1.0));

Only after resolving that the type of the aggregate is Complex can the compiler look inside and make sure Re and Im make sense.

This not-too-smart rule for aggregates helps prevent confusion on the part of developers reading the code. It also simplifies the compiler, and makes the overload resolution algorithm reasonably efficient.

11.5 Operator Overloading

We've seen *previously* (page 478) that we can define custom operators for any type. We've also seen that subprograms can be *overloaded* (page 487). Since operators are functions,

we're essentially talking about operator overloading, as we're defining the same operator (say + or -) for different types.

As another example of operator overloading, in the Ada standard library, operators are defined for the Complex type of the Ada.Numerics.Generic_Complex_Types package. This package contains not only the definition of the + operator for two objects of Complex type, but also for combination of Complex and other types. For instance, we can find these declarations:

This example shows that the + operator — as well as other operators — are being overloaded in the **Generic**_Complex_Types package.

```
1 In the Ada Reference Manual
```

- 6.6 Overloading of Operators²⁰⁶
- G.1.1 Complex Types²⁰⁷

11.6 Operator Overriding

We can also override operators of derived types. This allows for modifying the behavior of operators for the corresponding derived types.

To override an operator of a derived type, we simply implement a function for that operator. This is the same as how we implement custom operators (as we've seen previously).

As an example, when adding two fixed-point values, the result might be out of range, which causes an exception to be raised. A common strategy to avoid exceptions in this case is to saturate the resulting value. This strategy is typically employed in signal processing algorithms, for example.

In this example, we declare and use the 32-bit fixed-point type TQ31:

Listing 51: fixed point.ads

```
1 package Fixed_Point is
2
3 D : constant := 2.0 ** (-31);
4 type TQ31 is delta D range -1.0 .. 1.0 - D;
5
6 end Fixed Point;
```

Listing 52: show_sat_op.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Fixed_Point; use Fixed_Point;
procedure Show_Sat_Op is
A, B, C : TQ31;
begin
```

(continues on next page)

²⁰⁶ http://www.ada-auth.org/standards/22rm/html/RM-6-6.html
 ²⁰⁷ http://www.ada-auth.org/standards/22rm/html/RM-G-1-1.html

```
A := TQ31'Last;
7
      B := TQ31'Last;
8
      C := A + B;
9
10
      Put_Line (A'Image & " + "
11
                 & B'Image & " = "
12
                  & C'Image);
13
14
      A := TQ31'First;
15
      B := TQ31'First;
16
      C := A + B;
17
18
      Put_Line (A'Image & " + "
19
                 & B'Image & " = "
20
                  & C'Image);
21
22
   end Show_Sat_Op;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Operator_Overriding.Fixed_

→Point_Exception

MD5: 15d8860773ec7c0e505d0ee94781ae14
```

Runtime output

raised CONSTRAINT_ERROR : show_sat_op.adb:9 overflow check failed

Here, we're using the standard + operator, which raises a Constraint_Error exception in the C := A + B; statement due to an overflow. Let's now override the addition operator and enforce saturation when the result is out of range:

Listing 53: fixed point.ads

```
package Fixed_Point is
D : constant := 2.0 ** (-31);
type TQ31 is delta D range -1.0 .. 1.0 - D;
function "+" (Left, Right : TQ31)
return TQ31;
end Fixed_Point;
```



```
package body Fixed_Point is
1
2
      function "+" (Left, Right : TQ31)
3
                    return TQ31
4
      is
5
         type TQ31_2 is
6
7
           delta TQ31'Delta
            range TQ31'First * 2.0 .. TQ31'Last * 2.0;
8
9
            : constant TQ31_2 := TQ31_2 (Left);
         L
10
         R : constant TQ31_2 := TQ31_2 (Right);
11
         Res : TQ31_2;
12
      begin
13
```

(continues on next page)

```
Res := L + R;
14
15
          if Res > TQ31_2 (TQ31'Last) then
16
             return TQ31'Last;
17
          elsif Res < TQ31_2 (TQ31'First) then
18
             return TQ31'First;
19
          else
20
             return TQ31 (Res);
21
          end if;
22
       end "+";
23
24
   end Fixed_Point;
25
```

Listing 55: show_sat_op.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   with Fixed Point; use Fixed Point;
2
3
   procedure Show Sat Op is
4
      A, B, C : TQ31;
5
   begin
6
      A := TQ31'Last;
7
      B := T031'Last;
8
      C := A + B;
9
10
      Put Line (A'Image & " + "
11
                 & B'Image & " = "
12
                 & C'Image);
13
14
      A := TQ31'First;
15
      B := TQ31'First;
16
      C := A + B;
17
18
      Put Line (A'Image & " + "
19
                 & B'Image & " = "
20
                 & C'Image);
21
22
   end Show Sat Op;
23
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Operator_Overriding.Fixed_ ⇔Point_Operator_Overloading MD5: 6317bcf9c278c01f86dbdcb761d86240

Runtime output

0.9999999995 + 0.999999995 = 0.9999999995-1.0000000000 + -1.0000000000 = -1.0000000000

In the implementation of the overridden + operator of the TQ31 type, we declare another type (TQ31_2) with a wider range than TQ31. We use variables of the TQ31_2 type to perform the actual addition, and then we verify whether the result is still in TQ31's range. If it is, we simply convert the result *back* to the TQ31 type. Otherwise, we saturate it — using either the first or last value of the TQ31 type.

When overriding operators, the overridden operator replaces the original one. For example, in the A + B operation of the Show_Sat_Op procedure above, we're using the overridden version of the + operator, which performs saturation. Therefore, this operation doesn't raise an exception (as it was the case with the original + operator).

11.7 Nonreturning procedures

Usually, when calling a procedure P, we expect that it returns to the caller's *thread of control* after performing some action in the body of P. However, there are situations where a procedure never returns. We can indicate this fact by using the No_Return aspect in the subprogram declaration.

A typical example is that of a server that is designed to run forever until the process is killed or the machine where the server runs is switched off. This server can be implemented as an endless loop. For example:

Listing 56: servers.ads

```
package Servers is
procedure Run_Server
with No_Return;
end Servers:
```

Listing 57: servers.adb

```
package body Servers is
1
2
       procedure Run_Server is
3
       begin
4
          pragma Warnings
5
            (Off,
6
             "implied return after this statement");
7
          while True loop
8
             -- Processing happens here...
9
             null;
10
          end loop;
11
      end Run_Server;
12
13
   end Servers;
14
```

Listing 58: show_endless_loop.adb

```
with Servers; use Servers;
procedure Show_Endless_Loop is
begin
Run_Server;
end Show Endless Loop;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Nonreturning_Procedures.

→Server_Proc

MD5: 3f859b6e2aca8e31367658632e84126c
```

In this example, Run_Server doesn't exit from the **while True** loop, so it never returns to the Show_Endless_Loop procedure.

The same situation happens when we call a procedure that raises an exception unconditionally. In that case, exception handling is triggered, so that the procedure never returns to the caller. An example is that of a logging procedure that writes a message before raising an exception internally:

Listing 59: loggers.ads

```
1 package Loggers is
2
3 Logged_Failure : exception;
4
5 procedure Log_And_Raise (Msg : String)
6 with No_Return;
7
8 end Loggers;
```

Listing 60: loggers.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Loggers is
3
4
      procedure Log_And_Raise (Msg : String) is
5
      begin
6
         Put_Line (Msg);
7
          raise Logged_Failure;
8
      end Log_And_Raise;
9
10
  end Loggers;
11
```

Listing 61: show_no_return_exception.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
   with Loggers;
                    use Loggers;
2
3
   procedure Show No Return Exception is
4
      Check_Passed : constant Boolean := False;
5
   begin
6
      if not Check Passed then
7
         Log And Raise ("Check failed!");
8
         Put_Line ("This line will not be reached!");
9
      end if;
10
  end Show No Return Exception;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Nonreturning_Procedures.Log_ →Exception MD5: 10b4933d8c862d14ade54935cbd2b541

In this example, Log_And_Raise writes a message to the user and raises the Logged_Failure, so it never returns to the Show_No_Return_Exception procedure.

We could implement exception handling in the Show_No_Return_Exception procedure, so that the Logged_Failure exception could be handled there after it's raised in Log_And_Raise. However, this wouldn't be considered a *normal* return to the procedure because it wouldn't return to the point where it should (i.e. to the point where Put_Line is about to be called, right after the call to the Log_And_Raise procedure).

If a nonreturning procedure returns nevertheless, this is considered a program error, so that the Program_Error exception is raised. For example:

Listing 62: loggers.ads

package Loggers is

(continues on next page)

```
3 Logged_Failure : exception;
4
5 procedure Log_And_Raise (Msg : String)
6 with No_Return;
7
8 end Loggers;
```

Listing 63: loggers.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Loggers is
3
4
      procedure Log_And_Raise (Msg : String) is
5
      begin
6
         Put Line (Msg);
7
      end Log And Raise;
8
9
   end Loggers;
10
```

Listing 64: show_no_return_exception.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Loggers;
                    use Loggers;
3
   procedure Show_No_Return_Exception is
4
      Check_Passed : constant Boolean := False;
5
   begin
6
      if not Check Passed then
7
         Log And Raise ("Check failed!");
8
         Put Line ("This line will not be reached!");
9
      end if;
10
  end Show_No_Return_Exception;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Nonreturning_Procedures.

⇔Erroneous_Log_Exception

MD5: e44fd8df0529dda5749e85b9e300a999
```

Build output

Runtime output

Check failed!

raised PROGRAM_ERROR : loggers.adb:7 implicit return with No_Return

Here, Program_Error is raised when Log_And_Raise returns to the Show_No_Return_Exception procedure.

1 In the Ada Reference Manual

6.5.1 Nonreturning Subprograms²⁰⁸

11.8 Inline subprograms

Inlining²⁰⁹ refers to a kind of optimization where the code of a subprogram is expanded at the point of the call in place of the call itself.

In modern compilers, inlining depends on the optimization level selected by the user. For example, if we select the higher optimization level, the compiler will perform automatic inlining agressively.

1 In the GNAT toolchain

The highest optimization level (-03) of GNAT performs aggressive automatic inlining. This could mean that this level inlines too much rather than not enough. As a result, the cache may become an issue and the overall performance may be worse than the one we would achieve by compiling the same code with optimization level 2 (-02). Therefore, the general recommendation is to not *just* select -03 for the optimized version of an application, but instead compare it the optimized version built with -02.

It's important to highlight that the inlining we're referring above happens automatically, so the decision about which subprogram is inlined depends entirely on the compiler. However, in some cases, it's better to reduce the optimization level and perform manual inlining instead of automatic inlining. We do that by using the Inline aspect.

Let's look at this example:

Listing 65: float_arrays.ads

```
package Float_Arrays is
1
2
      type Float_Array is
3
        array (Positive range <>) of Float;
4
5
      function Average (Data : Float_Array)
6
                          return Float
7
         with Inline;
8
9
   end Float Arrays;
10
```

```
Listing 66: float_arrays.adb
```

```
package body Float Arrays is
1
2
      function Average (Data : Float Array)
3
                          return Float
4
5
      is
         Total : Float := 0.0;
6
      begin
7
          for Value of Data loop
8
             Total := Total + Value;
9
          end loop:
10
          return Total / Float (Data'Length);
11
      end Average;
12
```

(continues on next page)

²⁰⁸ http://www.ada-auth.org/standards/22rm/html/RM-6-5-1.html
 ²⁰⁹ https://en.wikipedia.org/wiki/Inline_expansion

```
14 end Float_Arrays;
```

13

Listing 67: compute_average.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Float Arrays; use Float Arrays;
3
   procedure Compute Average is
5
      Values
                     : constant Float Array :=
6
                         (10.0, 11.0, 12.0, 13.0);
7
      Average Value : Float;
8
   beain
9
      Average Value := Average (Values);
10
      Put Line ("Average = "
11
                 & Float'Image (Average Value));
12
   end Compute Average;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Inline_Subprograms.Inlining_

→Float_Arrays

MD5: 246bc11e8a969d69873f416f583f450e
```

Runtime output

Average = 1.15000E+01

When compiling this example, the compiler will most probably inline Average in the Compute_Average procedure. Note, however, that the Inline aspect is just a *recommendation* to the compiler. Sometimes, the compiler might not be able to follow this recommendation, so it won't inline the subprogram.

These are some examples of situations where the compiler might not be able to inline a subprogram:

- when the code is too large,
- when it's too complicated for example, when it involves exception handling —, or
- when it contains tasks, etc.

1 In the GNAT toolchain

In order to effectively use the Inline aspect, we need to set the optimization level to at least -01 and use the -gnatn switch, which instructs the compiler to take the Inline aspect into account.

In addition to the Inline aspect, in GNAT, we also have the (implementation-defined) Inline_Always aspect. In contrast to the former aspect, however, the Inline_Always aspect isn't primarily related to performance. Instead, it should be used when the functionality would be incorrect if inlining was not performed by the compiler. Examples of this are procedures that insert Assembly instructions that only make sense when the procedure is inlined, such as memory barriers.

Similar to the Inline aspect, there might be situations where a subprogram has the Inline_Always aspect, but the compiler is unable to inline it. In this case, we get a compilation error from GNAT.

Note that we can use the Inline aspect for generic subprograms as well. When we do this, we indicate to the compiler that we wish it inlines all instances of that generic subprogram.

```
1 In the Ada Reference Manual
```

6.3.2 Inline Expansion of Subprograms²¹⁰

11.9 Null Procedures

Null procedures are procedures that don't have any effect, as their body is empty. We declare a null procedure by simply writing **is null** in its declaration. For example:

Listing 68: null procs.ads

```
1 package Null_Procs is
2
3 procedure Do_Nothing (Msg : String) is null;
4
5 end Null_Procs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Null_Procedures.Null_Proc_1
MD5: a8a801e6c71d8177db61e4aa131b8832
```

As expected, calling a null procedure doesn't have any effect. For example:

Listing 69: show_null_proc.adb

```
with Null_Procs; use Null_Procs;
procedure Show_Null_Proc is
begin
Do_Nothing ("Hello");
end Show_Null_Proc;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Null_Procedures.Null_Proc_1
MD5: 274eed0b0952b9aa7e422933ece42d86
```

Null procedures are equivalent to implementing a procedure with a body that only contains **null**. Therefore, the Do_Nothing procedure above is equivalent to this:

Listing 70: null_procs.ads

```
package Null_Procs is
procedure Do_Nothing (Msg : String);
end Null_Procs;
```

Listing 71: null_procs.adb

```
package body Null_Procs is
procedure Do_Nothing (Msg : String) is
```

(continues on next page)

²¹⁰ http://www.ada-auth.org/standards/22rm/html/RM-6-3-2.html

```
4 begin
5 null;
6 end Do_Nothing;
7
8 end Null_Procs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Null_Procedures.Null_Proc_1
MD5: d0c9dc9265ebbaa9603681182dee1d92
```

11.9.1 Null procedures and overriding

We can use null procedures as a way to simulate interfaces for non-tagged types — similar to what actual interfaces do for tagged types. For example, we may start by declaring a type and null procedures that operate on that type. For example, let's model a very simple API:

Listing 72: simple_storage.ads

```
package Simple_Storage is
1
2
      type Storage_Model is null record;
3
4
      procedure Set (S : in out Storage_Model;
5
                                  String) is null;
                      V :
6
      procedure Display (S : Storage_Model) is null;
7
8
  end Simple Storage;
9
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Null_Procedures.Simple_

→Storage_Model

MD5: 553e78bc15dcec1302be4b5f484ac21f
```

Here, the API of the Storage_Model type consists of the Set and Display procedures. Naturally, we can use objects of the Storage_Model type in an application, but this won't have any effect:

Listing 73: show_null_proc.adb

```
with Ada.Text IO;
                         use Ada.Text I0;
1
   with Simple_Storage; use Simple_Storage;
2
3
   procedure Show_Null_Proc is
4
      S : Storage_Model;
5
   begin
6
      Put Line ("Setting 24...");
7
      Set (S, "24");
8
      Display (S);
9
  end Show_Null_Proc;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Null_Procedures.Simple_

→Storage_Model

MD5: 523b3e7239e683f2d879caa9139106ca
```

Runtime output

Setting 24...

By itself, the Storage_Model type is not very useful. However, we can derive other types from it and override the null procedures. Let's say we want to implement the Integer_Storage type to store an integer value:

```
Listing 74: simple storage.ads
```

```
package Simple_Storage is
1
2
      type Storage_Model is null record;
3
4
      procedure Set (S : in out Storage_Model;
5
                       V :
                                  String) is null;
6
      procedure Display (S : Storage_Model) is null;
7
8
      type Integer_Storage is private;
9
10
      procedure Set (S : in out Integer Storage;
11
                       V :
                                  String);
12
      procedure Display (S : Integer Storage);
13
14
   private
15
16
      type Integer_Storage is record
17
         V : Integer := 0;
18
      end record;
19
20
   end Simple_Storage;
21
```

Listing 75: simple_storage.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Storage is
3
4
      procedure Set (S : in out Integer_Storage;
5
                      V :
                                 String) is
6
      begin
7
         S.V := Integer'Value (V);
8
      end Set;
9
10
      procedure Display (S : Integer_Storage) is
11
      begin
12
         Put_Line ("Value: " & S.V'Image);
13
      end Display;
14
15
   end Simple_Storage;
16
```

Listing 76: show_null_proc.adb

```
with Ada.Text IO;
                         use Ada.Text I0;
1
  with Simple_Storage; use Simple_Storage;
2
3
   procedure Show_Null_Proc is
4
      S : Integer_Storage;
5
   begin
6
      Put_Line ("Setting 24...");
7
      Set (S, "24");
8
      Display (S);
9
  end Show_Null_Proc;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Subprograms.Null_Procedures.Simple_

⇔Storage_Model

MD5: 55d491d1ef72fb7be2bf0d2a212f335b
```

Runtime output

Setting 24... Value: 24

In this example, we can view Storage_Model as a sort of interface for derived non-tagged types, while the derived types — such as Integer_Storage — provide the actual implementation.

The section on *null records* (page 179) contains an extended example that makes use of null procedures.

1 In the Ada Reference Manual

• 6.7 Null Procedures²¹¹

²¹¹ http://www.ada-auth.org/standards/22rm/html/RM-6-7.html

CHAPTER TWELVE

EXCEPTIONS

12.1 Classification of Errors

When we talk about errors and erroneous behavior in Ada, we can classify them in one of the four categories:

- compilation errors i.e. errors that an Ada compiler must detect at compilation time;
- runtime errors i.e. errors that are detected by an Ada-based application using checks at runtime;
- bounded errors;
- erroneous execution.

In this section, we discuss each of these categories.

1 In the Ada Reference Manual

• 1.1.5 Classification of Errors²¹²

12.1.1 Compilation errors

In the category of compilation errors, the goal is to prevent compilers from accepting illegal programs. Here, any program that doesn't follow the rules described in the Ada Reference Manual is considered illegal. Those rules include not only simple syntax errors, but also more complicated semantic rules, such as the ones concerning *accessibility levels* (page 645) for access types.

Note that Ada — in contrast to many programming languages, which can be quite permissive — tries to prevent as many errors as possible at compilation time because of its focus on safety. However, even though a wide range of errors can be detected at compilation time, this doesn't mean that a legal Ada program is free of errors. Therefore, using methods such as static analysis or unit testing is important.

12.1.2 Runtime errors

When a rule cannot be verified at compilation time, a common strategy is to have the compiler insert runtime checks into the resulting application. We see details about these checks later on when we discuss *checks and exceptions* (page 513).

A typical example is an *overflow check* (page 519). Consider a calculation using variables: if this calculation leads to a result that isn't representable with the underlying data types, we cannot possibly store a value into a register or memory that can be considered correct — so we have to detect this situation. Unfortunately, because we're using variables, we

²¹² http://www.ada-auth.org/standards/12rm/html/RM-1-1-5.html

obviously cannot verify the result of the calculation at compilation time, so we have to verify it at runtime.

As we've mentioned before, Ada strives for detecting as many erroneous conditions as possible, while other programming language would allow errors such as overflow errors to remain undetected — which would likely lead the application to misbehave. Those checks raise an exception if an erroneous condition is detected, so the programmer has the means — and the responsibility — to catch that exception and handle the situation properly (Note, however, that some of the runtime checks can be deactivated. We will discuss this topic later on.)

12.1.3 Bounded errors

For certain kinds of errors, the compiler might not be able to detect the error — neither at compilation time, nor with checks at runtime. Such errors are called bounded errors because their possible effects are *bounded*. In fact, the Ada Reference Manual describes each bounded error and its possible effects — one of those effects is raising the Program_Error exception.

Just as an example, consider the bounded error described in section 13.9.1 Data Validity²¹³, paragraphs 9:

If the representation of a scalar object does not represent a value of the object's subtype (perhaps because the object was not initialized), the object is said to have an invalid representation. It is a bounded error to evaluate the value of such an object. If the error is detected, either Constraint_Error or Program_Error is raised. Otherwise, execution continues using the invalid representation. The rules of the language outside this subclause assume that all objects have valid representations.

Let's see a code example:

```
Listing 1: show_bounded_error.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Bounded Error is
3
      subtype Int_1_10 is
4
        Integer range 1 .. 10;
5
6
                 : Int 1 10;
      T1
7
      I1 Overlay : Integer
8
        with Address => I1'Address,
9
                         Import,
10
                          Volatile;
11
   begin
12
      I1_Overlay := 0;
13
       - -
14
          We use this overlay to write an invalid
15
          value to I1.
16
17
      Put Line ("I1 = " & I1'Image);
18
                             19
       -- Bounded error: value in
20
       -- Il is out of range.
21
22
      I1 := I1 + 1;
23
      - -
24
       -- Bounded error: using value
25
      -- in operation that is out of
26
```

(continues on next page)

²¹³ http://www.ada-auth.org/standards/12rm/html/RM-13-9-1.html

27	range.
28	
29	<pre>Put_Line ("I1 = " & I1'Image);</pre>
30	<pre>end Show_Bounded_Error;</pre>

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Classification_Of_Errors.

→Data_Validity_Bounded_Error

MD5: 770ebb7b6e0015e373e96c0dce250caa
```

Runtime output

 $\begin{array}{rrrr} I1 &= & 0\\ I1 &= & 1 \end{array}$

In this example, we simulate a missing initialization by using an overlay (I1_Overlay). As a consequence, I1 has an invalid value that is out of the allowed range of the Int_1_10 subtype. This situation causes two bounded errors:

- a bounded error when I1 is evaluated in the call to Image; and
- a bounded error when the value of the right-sided I1 is evaluated in the increment I1 := I1 + 1.

In the Ada Reference Manual

13.9.1 Data Validity²¹⁴

12.1.4 Erroneous execution

Erroneous execution is similar to bounded errors in the sense that having the compiler detect the erroneous condition at compilation time or at runtime isn't possible. However, unlike bounded errors, the effects are usually nondeterministic: a bound on possible effects is not described by the language.

Again, as an example of erroneous execution, consider the description from section 13.9.1 Data Validity²¹⁵, paragraph 12/3, which discusses the implications of using the Unchecked_Conversion function. Let's see a code example:

Listing 2: show_erroneous_execution.adb

```
with Ada.Text_IO; use Ada.Text_IO;
   with Ada.Unchecked Conversion;
2
   procedure Show Erroneous Execution is
4
      subtype Int 1 10 is
5
         Integer range 1 .. 10;
6
7
      function To Int 1 10 is new
8
        Ada.Unchecked Conversion
9
           (Source => Integer,
10
           Target => Int 1 10);
11
12
      I1 : Int_1_10 := To_Int_1_10 (0);
13
14
```

(continues on next page)

²¹⁴ http://www.ada-auth.org/standards/12rm/html/RM-13-9-1.html

²¹⁵ http://www.ada-auth.org/standards/12rm/html/RM-13-9-1.html

```
-- Bounded error
15
   begin
16
      Put_Line ("I1 = " & I1'Image);
17
18
      I1 := I1 + 1;
19
             ~~~~~
20
      - -
      -- Erroneous execution: using value
21
      -- in operation that is out of range.
22
23
      Put_Line ("I1 = " & I1'Image);
24
  end Show Erroneous Execution;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Classification_Of_Errors.

→Data_Validity_Erroneous_Execution

MD5: 19218e9bb2e153366dea9114a5e59314
```

Build output

Runtime output

 $\begin{array}{rrrr} I1 &= & 0\\ I1 &= & 1 \end{array}$

It is considered to be a bounded error to use the To_Int_1_10 function (based on Unchecked_Conversion) with a value that is invalid for the target data type. However, if we use the invalid value of I1 in an operation such as the I1 := I1 + 1 assignment, this leads to erroneous execution, and the effects are unpredictable: they aren't described in the Ada Reference Manual, as they are nondeterministic.

1 In the Ada Reference Manual

• 13.9.1 Data Validity²¹⁶

12.2 Asserts

When we want to indicate a condition in the code that must always be valid, we can use the pragma Assert. As the name implies, when we use this pragma, we're *asserting* some truth about the source-code. (We can also use the procedural form, as we'll see later.)

Important

Another method to assert the truth about the source-code is to use pre and post-conditions.

A simple assert has this form:

²¹⁶ http://www.ada-auth.org/standards/12rm/html/RM-13-9-1.html

Listing 3: show_pragma_assert.adb

```
1 procedure Show_Pragma_Assert is
2 I : constant Integer := 10;
3
4 pragma Assert (I = 10);
5 begin
6 null;
7 end Show_Pragma_Assert;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Asserts.Pragma_Assert_1
MD5: 8d40817304515169d0d5670904ca1e01
```

In this example, we're asserting that the value of I is always 10. We could also display a message if the assertion is false:

Listing 4: show_pragma_assert.adb

```
1 procedure Show_Pragma_Assert is
2 I : constant Integer := 11;
3
4 pragma Assert (I = 10, "I is not 10");
5 begin
6 null;
7 end Show_Pragma_Assert;
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Asserts.Pragma_Assert_2
MD5: b70fa67c92542ade39c388964ce12302

Build output

show_pragma_assert.adb:4:19: warning: assertion will fail at run time [-gnatw.a]

Runtime output

raised ADA.ASSERTIONS.ASSERTION_ERROR : I is not 10

Similarly, we can use the procedural form of Assert. For example, the code above can implemented as follows:

```
Listing 5: show_procedure_assert.adb
```

```
with Ada.Assertions; use Ada.Assertions;
procedure Show_Procedure_Assert is
I : constant Integer := 11;
begin
Assert (I = 10, "I is not 10");
end Show_Procedure_Assert;
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Asserts.Procedure_Assert
MD5: cbab23645ff89d4adffcaaddaeb6f0e3

Runtime output

raised ADA.ASSERTIONS.ASSERTION ERROR : I is not 10

Note that a call to Assert is simply translated to a check — and the Assertion_Error exception from the Ada.Assertions package being raised in the case that the check fails. For example, the code above roughly corresponds to this:

Listing 6: show_assertion_error.adb

```
with Ada.Assertions; use Ada.Assertions;
1
   procedure Show Assertion Error is
3
      I : constant Integer := 11;
4
5
   begin
6
      if I /= 10 then
7
         raise Assertion_Error with "I is not 10";
8
      end if;
9
10
  end Show Assertion Error;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Asserts.Assertion_Error MD5: 9c846acf998ca7adabd47c3b5a6ce39f

Runtime output

```
raised ADA.ASSERTIONS.ASSERTION_ERROR : I is not 10
```

1 In the Ada Reference Manual

11.4.2 Pragmas Assert and Assertion Policy²¹⁷

12.3 Assertion policies

We can activate and deactivate assertions based on assertion policies. We can do that by using the pragma Assertion_Policy. As an argument to this pragma, we indicate whether a specific policy must be checked or ignored.

For example, we can deactivate assertion checks by specifying Assert => Ignore. Similarly, we can activate assertion checks by specifying Assert => Check. Let's see a code example:

Listing 7: show_pragma_assertion_policy.adb

```
1 procedure Show_Pragma_Assertion_Policy is
2 I : constant Integer := 11;
3
4 pragma Assertion_Policy (Assert => Ignore);
5 begin
6 pragma Assert (I = 10);
7 end Show_Pragma_Assertion_Policy;
```

Code block metadata

²¹⁷ http://www.ada-auth.org/standards/22rm/html/RM-11-4-2.html

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Assertion_Policies.Pragma_ ⇔Assertion_Policy_1 MD5: 39b8aa4a34b6169c03b54074f4136519

Build output

Here, we're specifying that asserts shall be ignored. Therefore, the call to the pragma Assert doesn't raise an exception. If we replace Ignore with Check in the call to Assertion_Policy, the assert will raise the Assertion_Error exception.

The following table presents all policies that we can set:

Policy	Descripton
Assert	Check assertions
Static_Predicate	Check static predicates
Dynamic_Predicate	Check dynamic predicates
Pre	Check pre-conditions
Pre'Class	Check pre-condition of classes of tagged types
Post	Check post-conditions
Post'Class	Check post-condition of classes of tagged types
Type_Invariant	Check type invariants
Type_Invariant'Class	Check type invariant of classes of tagged types

1 In the GNAT toolchain

Compilers are free to include policies that go beyond the ones listed above. For example, GNAT includes the following policies — called *assertion kinds* in this context:

- Assertions
- Assert_And_Cut
- Assume
- Contract_Cases
- Debug
- Ghost
- Initial_Condition
- Invariant
- Invariant'Class
- Loop_Invariant
- Loop_Variant
- Postcondition
- Precondition
- Predicate
- Refined_Post
- Statement_Assertions
- Subprogram_Variant

Also, in addtion to Check and Ignore, GNAT allows you to set a policy to Disable and Suppressible.

You can read more about them in the GNAT Reference Manual²¹⁸.

You can specify multiple policies in a single call to Assertion_Policy. For example, you can activate all policies by writing:

Listing 8: show_multiple_assertion_policies.adb

```
procedure Show Multiple Assertion Policies is
1
      pragma Assertion_Policy
2
        (Assert
                               => Check,
3
         Static Predicate
                               => Check,
4
         Dynamic_Predicate
                               => Check,
5
         Pre
                               => Check,
6
         Pre'Class
                               => Check,
7
         Post
                               => Check,
8
         Post'Class
                              => Check,
9
         Type_Invariant => Check,
10
         Type_Invariant'Class => Check);
11
   begin
12
13
      null:
   end Show_Multiple_Assertion_Policies;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Assertion_Policies.Multiple_

⇔Assertion_Policies

MD5: 3abbf97160b755b84cc4f7e652ca5551
```

1 In the GNAT toolchain

With GNAT, policies can be specified in multiple ways. In addition to calls to Assertion_Policy, you can use configuration pragmas files²¹⁹. You can use these files to specify all pragmas that are relevant to your application, including Assertion_Policy. In addition, you can manage the granularity for those pragmas. For example, you can use a global configuration pragmas file for your complete application, or even different files for each source-code file you have.

Also, by default, all policies listed in the table above are deactivated, i.e. they're all set to Ignore. You can use the command-line switch -gnata to activate them.

Note that the Assert procedure raises an exception independently of the assertion policy (Assert => Ignore)). For example:

Listing 9: show_assert_procedure_policy.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Assertions; use Ada.Assertions;
procedure Show_Assert_Procedure_Policy is
pragma Assertion_Policy (Assert => Ignore);
I : constant Integer := 1;
begin
```

(continues on next page)

²¹⁸ https://gcc.gnu.org/onlinedocs/gnat rm/Pragma-Assertion 005fPolicy

²¹⁹ https://gcc.gnu.org/onlinedocs/gnat_ugn/The-Configuration-Pragmas-Files#The-Configuration-Pragmas-Files

```
9 Put_Line ("----- Pragma Assert -----");

10 pragma Assert (I = 0);

11

12 Put_Line ("---- Procedure Assert ----");

13 Assert (I = 0);

14

15 Put_Line ("Finished.");

16 end Show_Assert_Procedure_Policy;
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Assertion_Policies.Assert_ ⊶Procedure_Policy MD5: 7be3bab24d856081afeddabe40afc84f

Build output

Runtime output

```
----- Pragma Assert -----
---- Procedure Assert ----
```

raised ADA.ASSERTIONS.ASSERTION_ERROR : a-assert.adb:42

Here, the **pragma** Assert is ignored due to the assertion policy. However, the call to Assert is not ignored.

In the Ada Reference Manual

• 11.4.2 Pragmas Assert and Assertion_Policy²²⁰

12.4 Checks and exceptions

This table shows all language-defined checks and the associated exceptions:

Check	Exception
Access_Check	Constraint_Error
Discriminant_Check	Constraint_Error
Division_Check	Constraint_Error
Index_Check	Constraint_Error
Length_Check	Constraint_Error
Overflow_Check	Constraint_Error
Range_Check	Constraint_Error
Tag_Check	Constraint_Error
Accessibility_Check	Program_Error
Allocation_Check	Program_Error
Elaboration_Check	Program_Error
Program_Error_Check	Program_Error
Storage_Check	Storage_Error
Tasking_Check	Tasking_Error

²²⁰ http://www.ada-auth.org/standards/22rm/html/RM-11-4-2.html

In addition, we can use All_Checks to refer to all those checks above at once.

Let's discuss each check and see code examples where those checks are performed. Note that all examples are erroneous, so please avoid reusing them elsewhere.

12.4.1 Access Check

As you know, an object of an access type might be null. It would be an error to dereference this object, as it doesn't indicate a valid position in memory. Therefore, the access check verifies that an access object is not null when dereferencing it. For example:

Listing 10: show_access_check.adb

```
1 procedure Show_Access_Check is
2
3 type Integer_Access is access Integer;
4
5 AI : Integer_Access;
6 begin
7 AI.all := 10;
8 end Show_Access_Check;
```

Code block metadata

Build output

Runtime output

raised CONSTRAINT_ERROR : show_access_check.adb:7 access check failed

Here, the value of AI is null by default, so we cannot dereference it.

The access check also performs this verification when assigning to a subtype that excludes null (**not null access**). (You can find more information about this topic in the section about *not null access* (page 664).) For example:

```
Listing 11: show_access_check.adb
```

```
procedure Show_Access_Check is
1
2
      type Integer Access is
3
        access all Integer;
4
5
      type Safe_Integer_Access is
6
        not null access all Integer;
7
8
      AI : Integer_Access;
9
      SAI : Safe_Integer_Access := new Integer;
10
11
   begin
12
```

(continues on next page)

```
13 SAI := Safe_Integer_Access (AI);
14 end Show Access Check;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.Access_

→Check_2

MD5: 47895a404e2a111476cd67f43c12d4b5
```

Build output

Runtime output

```
raised CONSTRAINT_ERROR : show_access_check.adb:13 access check failed
```

Here, the value of AI is null (by default), so we cannot assign it to SAI because its type excludes null.

Note that, if we remove the := **new Integer** assignment from the declaration of SAI, the null exclusion fails in the declaration itself (because the default value of the access type is **null**).

12.4.2 Discriminant Check

As we've seen earlier, a variant record is a record with discriminants that allows for changing its structure. In operations such as an assignment, it's important to ensure that the discriminants of the objects match — i.e. to ensure that the structure of the objects matches. The discriminant check verifies whether this is the case. For example:

Listing 12: show_discriminant_check.adb

```
procedure Show_Discriminant_Check is
1
2
       type Rec (Valid : Boolean) is record
3
          case Valid is
4
             when True =>
5
                Counter : Integer;
6
             when False =>
7
                null;
8
          end case;
9
       end record;
10
11
       R : Rec (Valid => False);
12
   begin
13
       R := (Valid => True,
14
             Counter \Rightarrow 10);
15
   end Show_Discriminant_Check;
16
```

Code block metadata

Build output

Runtime output

raised CONSTRAINT_ERROR : show_discriminant_check.adb:14 discriminant check failed

Here, R's discriminant (Valid) is **False**, so we cannot assign an object whose Valid discriminant is **True**.

Also, when accessing a component, the discriminant check ensures that this component exists for the current discriminant value:

```
Listing 13: show_discriminant_check.adb
```

```
procedure Show Discriminant Check is
1
2
      type Rec (Valid : Boolean) is record
3
          case Valid is
4
             when True =>
5
                Counter : Integer;
6
             when False =>
7
                null;
8
          end case;
9
      end record;
10
11
      R : Rec (Valid => False);
12
      I : Integer;
13
   begin
14
      I := R.Counter;
15
   end Show_Discriminant_Check;
16
```

Code block metadata

Build output

Runtime output

raised CONSTRAINT_ERROR : show_discriminant_check.adb:15 discriminant check failed

Here, R's discriminant (Valid) is **False**, so we cannot access the Counter component, for it only exists when the Valid discriminant is **True**.

12.4.3 Division Check

The division check verifies that we're not trying to divide a value by zero when using the /, rem and mod operators. For example:

```
Listing 14: ops.ads
```

```
package Ops is
1
       function Div_Op (A, B : Integer)
2
                         return Integer is
3
         (A / B);
4
5
      function Rem Op (A, B : Integer)
6
                         return Integer is
7
         (A rem B);
8
9
       function Mod_Op (A, B : Integer)
10
                         return Integer is
11
         (A mod B):
12
   end Ops;
13
```

Listing 15: show_division_check.adb

```
with Ops; use Ops;
1
2
  procedure Show_Division_Check is
3
      I : Integer;
4
  begin
5
      I := Div_Op (10, 0);
6
      I := Rem_Op (10, 0);
7
      I := Mod_Op (10, 0);
8
  end Show_Division_Check;
9
```

Code block metadata

Runtime output

raised CONSTRAINT_ERROR : ops.ads:4 divide by zero

All three calls in the Show_Division_Check procedure — to the Div_Op, Rem_Op and Mod_Op functions — can raise an exception because we're using 0 as the second argument, which makes the division check in those functions fail.

12.4.4 Index Check

We use indices to access components of an array. An index check verifies that the index we're using to access a specific component is within the array's bounds. For example:

Listing 16: show_index_check.adb

```
1 procedure Show_Index_Check is
2
3 type Integer_Array is
4 array (Positive range <>) of Integer;
5
6 function Value_Of (A : Integer_Array;
```

(continues on next page)

```
I : Integer)
7
                            return Integer
8
       is
9
          type Half_Integer_Array is new
10
            Integer_Array (A'First ..
11
                            A'First + A'Length / 2);
12
13
          A_2 : Half_Integer_Array := (others => 0);
14
       begin
15
          return A_2 (I);
16
      end Value_0f;
17
18
19
      Arr_1 : Integer_Array (1 .. 10) :=
                  (others => 1);
20
21
   begin
22
      Arr_1 (10) := Value_Of (Arr_1, 10);
23
24
   end Show_Index_Check;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.Index_

→Check

MD5: fa791718701c4ac805badf368df9064e
```

Runtime output

```
raised CONSTRAINT_ERROR : show_index_check.adb:16 index check failed
```

The range of A_2 — which is passed as an argument to the Value_Of function — is 1 to 6. However, in that function call, we're trying to access position 10, which is outside A_2 's bounds.

12.4.5 Length Check

In array assignments, both arrays must have the same length. To ensure that this is the case, a length check is performed. For example:

```
Listing 17: show_length_check.adb
```

```
procedure Show_Length_Check is
1
2
      type Integer Array is
3
        array (Positive range <>) of Integer;
4
5
      procedure Assign (To : out Integer Array;
6
                         From :
                                    Integer Array) is
7
      begin
8
         To := From;
9
      end Assign;
10
11
      Arr_1 : Integer_Array (1 .. 10);
12
      Arr_2 : Integer_Array (1 .. 9) :=
13
                 (others => 1);
14
15
   begin
16
      Assign (Arr_1, Arr_2);
17
   end Show Length Check;
18
```

Code block metadata

Runtime output

raised CONSTRAINT_ERROR : show_length_check.adb:9 length check failed

Here, the length of Arr_1 is 10, while the length of Arr_2 is 9, so we cannot assign Arr_2 (From parameter) to Arr_1 (To parameter) in the Assign procedure.

12.4.6 Overflow Check

Operations on scalar objects might lead to overflow, which, if not checked, lead to wrong information being computed and stored. Therefore, an overflow check verifies that the value of a scalar object is within the base range of its type. For example:

Listing 18: show_overflow_check.adb

```
1 procedure Show_Overflow_Check is
2 A, B : Integer;
3 begin
4 A := Integer'Last;
5 B := 1;
6
7 A := A + B;
8 end Show_Overflow_Check;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.

⊶Overflow_Check

MD5: baa46d9085cbd14863aaa7e24dc7b9cc
```

Build output

Runtime output

raised CONSTRAINT_ERROR : show_overflow_check.adb:7 overflow check failed

In this example, A already has the last possible value of the **Integer**'Base range, so increasing it by one causes an overflow error.

12.4.7 Range Check

The range check verifies that a scalar value is within a specific range — for instance, the range of a subtype. Let's see an example:

Listing 19: show_range_check.adb

```
procedure Show_Range_Check is
subtype Int_1_10 is Integer range 1 .. 10;
I : Int_1_10;
begin
I := 11;
end Show_Range_Check;
```

Code block metadata

Build output

Runtime output

raised CONSTRAINT_ERROR : show_range_check.adb:8 range check failed

In this example, we're trying to assign 11 to the variable I of the Int_1_10 subtype, which has a range from 1 to 10. Since 11 is outside that range, the range check fails.

12.4.8 Tag Check

The tag check ensures that the tag of a tagged object matches the expected tag in a dispatching operation. For example:

Listing 20: p.ads

```
package P is
type T is tagged null record;
type T1 is new T with null record;
type T2 is new T with null record;
end P;
```

Listing 21: show_tag_check.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with Ada.Tags;
2
3
   with P;
                      use P;
4
5
   procedure Show Tag Check is
6
7
      A1 : T'Class := T1'(null record);
8
      A2 : T'Class := T2'(null record);
9
10
   begin
11
```

(continues on next page)

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.Tag_

→Check

MD5: 5a685be7804200a884649f54c175ee42
```

Runtime output

```
A1'Tag: P.T1
A2'Tag: P.T2
raised CONSTRAINT_ERROR : show_tag_check.adb:17 tag check failed
```

Here, A1 and A2 have different tags:

- A1'Tag = T1'Tag, while
- A2'Tag = T2'Tag.

Since the tags don't match, the tag check fails in the assignment of A1 to A2.

12.4.9 Accessibility Check

The accessibility check verifies that the accessibility level of an entity matches the expected level. We discuss accessibility levels *in a later chapter* (page 645).

Let's look at an example that mixes access types and anonymous access types. Here, we use an anonymous access type in the declaration of A1 and a named access type in the declaration of A2:

Listing 22: p.ads

```
package P is
type T is tagged null record;
type T_Class is access all T'Class;
end P;
```

Listing 23: show_accessibility_check.adb

```
with P; use P;
1
2
   procedure Show Accessibility Check is
3
4
      A1 : access T'Class := new T:
5
      A2 : T Class;
6
7
   begin
8
      A2 := T Class (A1);
9
10
   end Show Accessibility Check;
11
```

Code block metadata

Build output

Runtime output

```
raised PROGRAM_ERROR : show_accessibility_check.adb:9 accessibility check failed
```

The anonymous type (access T'Class), which is used in the declaration of A1, doesn't have the same accessibility level as the T_Class type. Therefore, the accessibility check fails during the T_Class (A1) conversion.

We can see the accessibility check failing in this example as well:

Listing 24: show_accessibility_check.adb

```
with P; use P;
1
   procedure Show_Accessibility_Check is
3
      A : access T'Class := new T;
5
6
      procedure P (A : T_Class) is null;
7
8
   begin
9
      P (T_Class (A));
10
11
   end Show Accessibility Check;
12
```

Code block metadata

Build output

Runtime output

raised PROGRAM_ERROR : show_accessibility_check.adb:10 accessibility check failed

Again, the check fails in the T_Class (A) conversion and raises a Program_Error exception.

12.4.10 Allocation Check

The allocation check ensures, when a task is about to be created, that its master has not been completed. Also, it ensures that the finalization has not started.

This is an example adapted from AI-00280²²¹:

```
Listing 25: p.ads
```

```
with Ada.Finalization;
1
   with Ada.Unchecked_Deallocation;
2
3
   package P is
4
      type T1 is new
5
        Ada.Finalization.Controlled with null record;
6
      procedure Finalize (X : in out T1);
7
8
9
      type T2 is new
        Ada.Finalization.Controlled with null record;
10
      procedure Finalize (X : in out T2);
11
12
      X1 : T1;
13
14
      type T2_Ref is access T2;
15
      procedure Free is new
16
        Ada.Unchecked_Deallocation (T2, T2_Ref);
17
   end P;
18
```

```
Listing 26: p.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body P is
3
4
      procedure Finalize (X : in out T1) is
5
         X2 : T2_Ref := new T2;
6
7
      begin
          Put_Line ("Finalizing T1...");
8
          Free (X2);
9
      end Finalize;
10
11
      procedure Finalize (X : in out T2) is
12
      begin
13
          Put Line ("Finalizing T2...");
14
      end Finalize;
15
16
   end P;
17
```

²²¹ http://www.ada-auth.org/cgi-bin/cvsweb.cgi/ais/ai-00280.txt?rev=1.12&raw=N

Listing 27: show_allocation_check.adb

```
with P; use P;
procedure Show_Allocation_Check is
X2 : T2_Ref := new T2;
begin
Free (X2);
end Show_Allocation_Check;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.

→Allocation_Check

MD5: 915e8ab21e550c981503c014bcceade1
```

Runtime output

```
Finalizing T2...
```

raised PROGRAM_ERROR : finalize/adjust raised exception

Here, in the finalization of the X1 object of T1 type, we're trying to create an object of T2 type while the finalization of the master has already started. (Note that X1 was declared in the P package.) This is forbidden, so the allocation check raises a Program_Error exception.

12.4.11 Elaboration Check

The elaboration check verifies that subprograms — or protected entries, or task activations — have been elaborated before being called.

This is an example adapted from AI-00064²²²:

Listing 28: p.ads

1 function P return Integer;

Listing 29: p.adb

```
1 function P return Integer is
2 begin
3 return 1;
4 end P;
```

Listing 30: show_elaboration_check.adb

```
with P:
1
2
   procedure Show Elaboration Check is
3
4
      function F return Integer;
5
6
      type Pointer_To_Func is
7
        access function return Integer;
8
9
      X : constant Pointer_To_Func := P'Access;
10
11
      Y : constant Integer := F;
12
```

(continues on next page)

 $^{222}\ http://www.ada-auth.org/cgi-bin/cvsweb.cgi/ais/ai-00064.txt?rev=1.12\&raw=N$

```
I3 Z : constant Pointer_To_Func := X;
I4
I5 -- Renaming-as-body
I6 function F return Integer renames Z.all;
I7 begin
I8 null;
I9 end Show_Elaboration_Check;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.

→Elaboration_Check

MD5: 80a39df912aae8788296f81ee9d4a79e
```

Build output

```
show_elaboration_check.adb:12:28: warning: cannot call "F" before body seen_

→[enabled by default]

show_elaboration_check.adb:12:28: warning: Program_Error will be raised at run_

→time [enabled by default]
```

Runtime output

raised PROGRAM_ERROR : show_elaboration_check.adb:12 access before elaboration

This is a curious example: first, we declare a function F and assign the value returned by this function to constant Y in its declaration. Then, we declare F as a renamed function, thereby providing a body to F — this is called renaming-as-body. Consequently, the compiler doesn't complain that a body is missing for function F. (If you comment out the function renaming, you'll see that the compiler can then detect the missing body.) Therefore, at runtime, the elaboration check fails because the body of the first declaration of the F function is actually missing.

12.4.12 Program_Error_Check

\rm 1 Note

This concept was introduced in Ada 2022.

As we've seen before, there are three checks that may raise a Program_Error exception: the Accessibility_Check, the Allocation_Check and the Elaboration_Check. In addition to that, we have the Program_Error_Check, which is actually a collection of various different checks that may raise a Program_Error, but don't have a category for themselves.

For completeness, these are the error conditions checked by the Program_Error_Check (listed in the Action Item (AI) 12-0309 document²²³), according to their definition in the Ada Reference Manual:

²²³ http://www.ada-auth.org/cgi-bin/cvsweb.cgi/ai12s/ai12-0309-1.txt?rev=1.5&raw=N

Ada Refer- ence Manual	Para- Description grapl
3.2.4 Subtype Predicates ²²⁴	(29.1 It checks that subtypes with predicates are not used to index an array in generic units.
5.5 Loop State- ments ²²⁵	(8.1/! It checks that the maximum number of chunks for statement-level parallelism is greater than zero.
6.4.1 Parameter Associations ²²⁶	(13.4 It checks that the default value of an out parameter is convert- ible: an error occurs when we have an out parameter with De- fault_Value, and the actual is a view conversion of an unrelated type that does not have Default Value.
12.5.1 For- mal Private and Derived Types ²²⁷	(23.3 It checks that there is no misuse of functions in a generic with a class-wide actual type.
13.3 Op- erational and Repre- sentation Attributes ²²⁸	(75.1 It checks that there are no colliding External_Tag values.
B.3.3 Unchecked Union Types ²²⁹	(22/2 It checks that there is no misuse of operations of Unchecked_Unions without inferable discriminants.

In the Ada Reference Manual
 11.5 Suppressing Checks²³⁰
• 3.2.4 Subtype Predicates ²³¹
• 5.5 Loop Statements ²³²
 6.4.1 Parameter Associations²³³
 12.5.1 Formal Private and Derived Types²³⁴
 13.3 Operational and Representation Attributes²³⁵
B.3.3 Unchecked Union Types ²³⁶

<sup>http://www.ada-auth.org/standards/22rm/html/RM-3-2-4.html
http://www.ada-auth.org/standards/22rm/html/RM-5-5.html
http://www.ada-auth.org/standards/22rm/html/RM-6-4-1.html
http://www.ada-auth.org/standards/22rm/html/RM-12-5-1.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-15-5.html
http://www.ada-auth.org/standards/22rm/html/RM-15-5.html
http://www.ada-auth.org/standards/22rm/html/RM-15-5.html
http://www.ada-auth.org/standards/22rm/html/RM-15-5.html
http://www.ada-auth.org/standards/22rm/html/RM-5-5.html
http://www.ada-auth.org/standards/22rm/html/RM-5-5.html
http://www.ada-auth.org/standards/22rm/html/RM-12-5-1.html
http://www.ada-auth.org/standards/22rm/html/RM-12-5-1.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
http://www.ada-auth.org/standards/22rm/html/RM-13-3.html</sup>

Example of a Program_Error_Check

Just as an example, let's look at the check for subtype predicates in generic units:

```
Listing 31: some_generic_package.ads
```

```
1 generic
2 type R is (<>);
3 package Some_Generic_Package is
4 procedure Process;
5 end Some_Generic_Package;
```

Listing 32: some_generic_package.adb

```
package body Some_Generic_Package is
1
2
      procedure Process is
3
          type Arr is
4
5
            array (R) of Integer;
6
          Dummy : Arr := (others => 0);
7
      begin
8
          null;
9
      end Process;
10
11
   end Some_Generic_Package;
12
```

Listing 33: show_subtype_predicate_programm_error.adb

```
with Some Generic Package;
1
2
   procedure Show Subtype Predicate Programm Error is
3
4
      type Custom_Range is range 1 .. 5
5
        with Dynamic_Predicate =>
6
                4 not in Custom Range;
7
8
      package P is new
9
         Some Generic Package (Custom Range);
10
      use P;
11
   begin
12
      Process:
13
   end Show Subtype Predicate Programm Error;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.

→Subtype_Predicate_Programm_Error

MD5: b1a5cc579393162dedecb6b65b75eef4
```

Build output

Runtime output

raised PROGRAM_ERROR : some_generic_package.adb:5 improper use of generic subtype. →with predicate

Here, we're using the Custom_Range type for the formal type R in the instantiation of the generic package Some_Generic_Package. Since we use R as an index for the array type Arr (in the procedure Process), we cannot map a type to R that uses a predicate. Therefore, because Custom_Range type has a dynamic predicate, the Program_Error exception is raised.

12.4.13 Storage Check

The storage check ensures that the storage pool has enough space when allocating memory. Let's revisit an example that we *discussed earlier* (page 87):

```
Listing 34: custom_types.ads
```

```
package Custom_Types is
type UInt_7 is range 0 .. 127;
type UInt_7_Reserved_Access is access UInt_7
with Storage_Size => 8;
end Custom_Types;
```

Listing 35: show_storage_check.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Custom_Types; use Custom_Types;
3
4
   procedure Show_Storage_Check is
5
6
      RAV1, RAV2 : UInt_7_Reserved_Access;
7
8
   begin
9
      Put_Line ("Allocating RAV1...");
10
      RAV1 := new UInt_7;
11
12
      Put_Line ("Allocating RAV2...");
13
      RAV2 := new UInt_7;
14
15
      New Line;
16
  end Show Storage Check;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Checks_And_Exceptions.

→Storage_Check

MD5: 4e4bd284adb1c1d97f8f7563068c18de
```

Runtime output

```
Allocating RAV1...
Allocating RAV2...
raised STORAGE_ERROR : s-poosiz.adb:108 explicit raise
```

On each allocation (new UInt_7), a storage check is performed. Because there isn't enough

reserved storage space before the second allocation, the checks fails and raises a Storage_Error exception.

12.4.14 Tasking_Check

The **Tasking**_Check ensures that all tasks have been activated successfully and that no terminated task is called. If the check fails, a **Tasking**_Error exception is raised.

\rm 1 Note

This concept was introduced in Ada 2022. It was created to group all checks that might raise the **Task**ing_Error exception.

Let's look at a simple example:

Listing 36: workers.ads

```
1 package Workers is
2
3 task type Worker is
4 entry Start;
5 end Worker;
6
7 end Workers;
```

```
Listing 37: workers.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Workers is
3
4
       task body Worker is
5
       begin
6
           Put_Line ("Task has started.");
7
           delay 1.0;
8
           Put_Line ("Task has finished.");
9
       end Worker;
10
11
   end Workers;
12
```

Listing 38: show_tasking_check_error.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   with Workers;
                    use Workers;
2
3
   procedure Show_Tasking_Check_Error is
4
       W : Worker;
5
   begin
6
       Put Line ("W.Start...");
7
       W.Start;
8
       Put_Line ("Finished");
9
  end Show Tasking Check Error;
10
```

Code block metadata

Build output

workers.adb:5:05: warning: no accept for entry "Start" [enabled by default]

Runtime output

Task has started. W.Start... Task has finished.

raised TASKING_ERROR

In this example, the body of Worker doesn't have an **accept**. Therefore, no rendezvous can happen for the W.Start call. Since the task eventually terminates (as you can see in the user messages), the call to Start constitutes a call to a terminated task. This condition is checked by the **Tasking_Check**, which fails in this case, thereby raising a **Tasking_Error**.

12.5 Ada.Exceptions package

\rm Note

Parts of this section were originally published as Gem #142 : Exception-ally²³⁷

The standard Ada run-time library provides the package Ada.Exceptions. This package provides a number of services to help analyze exceptions.

Each exception is associated with a (short) message that can be set by the code that raises the exception, as in the following code:

```
raise Constraint_Error with "some message";
```

Historically

Since Ada 2005, we can use the **raise** Constraint_Error **with** "some message" syntax. In Ada 95, you had to call the Raise_Exception procedure:

Ada.Exceptions.Raise_Exception -- Ada 95
 (Constraint_Error'Identity, "some message");

In Ada 83, there was no way to do it at all.

The new syntax is now very convenient, and developers should be encouraged to provide as much information as possible along with the exception.

1 In the GNAT toolchain

The length of the message is limited to 200 characters by default in GNAT, and messages longer than that will be truncated.

In the Ada Reference Manual

11.4.1 The Package Exceptions²³⁸

```
<sup>237</sup> https://www.adacore.com/gems/gem-142-exceptions
```

```
<sup>238</sup> http://www.ada-auth.org/standards/22rm/html/RM-11-4-1.html
```

12.5.1 Retrieving exception information

Exceptions also embed information set by the run-time itself that can be retrieved by calling the Exception_Information function. The function Exception_Information also displays the Exception_Message.

```
For example:
```

```
exception
when E : others =>
Put_Line
(Ada.Exceptions.Exception_Information (E));
```

1 In the GNAT toolchain

In the case of GNAT, the information provided by an exception might include the source location where the exception was raised and a nonsymbolic traceback.

You can also retrieve this information individually. Here, you can use:

- the Exception_Name functions and its derivatives Wide_Exception_Name and Wide_Wide_Exception_Name to retrieve the name of an exception.
- the Exception_Message function to retrieve the message associated with an exception.

Let's see a complete example:

```
Listing 39: show_exception_info.adb
```

```
with Ada.Text IO;
                          use Ada.Text I0;
1
   with Ada.Exceptions; use Ada.Exceptions;
2
3
   procedure Show Exception Info is
4
5
      Custom_Exception : exception;
6
7
      procedure Nested is
8
      begin
9
          raise Custom_Exception
10
           with "We got a problem";
11
      end Nested;
12
13
   begin
14
      Nested;
15
16
   exception
17
      when E : others =>
18
          Put_Line ("Exception info: "
19
                     & Exception_Information (E));
20
          Put_Line ("Exception name:
21
                     & Exception Name (E));
22
          Put_Line ("Exception msg:
23
                     & Exception Message (E));
24
   end Show Exception Info;
25
```

12.5.2 Collecting exceptions

Save_Occurrence

You can save an exception occurrence using the Save_Occurrence procedure. (Note that a Save_Occurrence function exists as well.)

For example, the following application collects exceptions into a list and displays them after running the Test_Exceptions procedure:

Listing 40: exception_tests.ads

```
with Ada.Exceptions; use Ada.Exceptions;
1
2
   package Exception Tests is
3
4
      Custom Exception : exception;
5
6
      type All_Exception_Occur is
7
        array (Positive range <>) of
8
9
          Exception_Occurrence;
10
      procedure Test Exceptions
11
         (All_Occur : in out All_Exception_Occur;
12
         Last Occur :
                        out Integer);
13
14
   end Exception_Tests;
15
```

```
Listing 41: exception_tests.adb
```

```
package body Exception_Tests is
1
2
      procedure Save_To_List
3
                               Exception Occurrence;
         (E
                     .
4
         All_Occur : in out All_Exception_Occur:
5
         Last_Occur : in out Integer)
6
      is
7
8
          L : Integer renames Last_Occur;
9
          0 : All_Exception_Occur renames All_Occur;
10
      begin
         L := L + 1;
11
         if L > 0'Last then
12
             raise Constraint_Error
13
               with "Cannot save occurrence";
14
         end if;
15
16
          Save Occurrence (Target => 0 (L),
17
                            Source => E);
18
      end Save_To_List;
19
20
      procedure Test_Exceptions
21
         (All_Occur : in out All_Exception_Occur;
22
          Last_Occur : out Integer)
23
      is
24
25
          procedure Nested_1 is
26
          begin
27
             raise Custom Exception
28
               with "We got a problem";
29
          exception
30
             when E : others =>
31
                Save_To_List (E,
32
                               All_Occur,
33
```

(continues on next page)

```
Last_Occur);
34
          end Nested_1;
35
36
          procedure Nested_2 is
37
38
          begin
              raise Constraint_Error
39
                with "Constraint is not correct";
40
          exception
41
              when E : others =>
42
                 Save_To_List (E,
43
                                 All Occur,
44
                                 Last_Occur);
45
          end Nested_2;
46
47
48
       begin
          Last_Occur := 0;
49
50
          Nested_1;
51
          Nested 2;
52
       end Test_Exceptions;
53
54
   end Exception_Tests;
55
```

Listing 42: show_exception_info.adb

```
with Ada.Text IO;
                          use Ada.Text I0;
1
   with Ada.Exceptions; use Ada.Exceptions;
2
3
   with Exception_Tests; use Exception_Tests;
4
5
   procedure Show_Exception_Info is
6
      L : Integer;
7
      0 : All_Exception_Occur (1 .. 10);
8
   begin
9
      Test_Exceptions (0, L);
10
11
      for I in O 'First .. L loop
12
         Put Line (Exception Information (0 (I)));
13
      end loop;
14
   end Show_Exception_Info;
15
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Exceptions_Package.Save_ Goccurrence MD5: da0cc5db7039e1458dbcf8be49db969d

Runtime output

raised EXCEPTION_TESTS.CUSTOM_EXCEPTION : We got a problem

```
raised CONSTRAINT_ERROR : Constraint is not correct
```

In the Save_To_List procedure of the Exception_Tests package, we call the Save_Occurrence procedure to store the exception occurrence to the All_Occur array. In the Show_Exception_Info, we display all the exception occurrences that we collected.

Read and Write attributes

Similarly, we can use files to read and write exception occurrences. To do that, we can simply use the Read and Write attributes.

Listing 43:	exception	occurrence	_stream.adb

```
with Ada.Text_I0;
1
2
   with Ada.Streams.Stream IO;
3
   use Ada.Streams.Stream_I0;
4
5
   with Ada.Exceptions;
6
   use Ada.Exceptions;
7
8
   procedure Exception_Occurrence_Stream is
9
10
      Custom_Exception : exception;
11
12
      S : Stream Access;
13
14
      procedure Nested_1 is
15
      begin
16
          raise Custom_Exception
17
            with "We got a problem";
18
      exception
19
          when E : others =>
20
             Exception_Occurrence'Write (S, E);
21
      end Nested_1;
22
23
      procedure Nested 2 is
24
      begin
25
          raise Constraint Error
26
            with "Constraint is not correct";
27
       exception
28
          when E : others =>
29
             Exception_Occurrence'Write (S, E);
30
      end Nested_2;
31
32
      F
                  : File_Type;
33
       File_Name : constant String :=
34
                      "exceptions_file.bin";
35
   begin
36
      Create (F, Out File, File Name);
37
      S := Stream (F);
38
39
      Nested 1;
40
41
      Nested_2;
42
      Close (F);
43
44
      Read_Exceptions : declare
45
          E : Exception_Occurrence;
46
       begin
47
          Open (F, In File, File Name);
48
          S := Stream (F);
49
50
          while not End_Of_File (F) loop
51
             Exception_Occurrence'Read (S, E);
52
53
             Ada.Text_I0.Put_Line
54
                (Exception_Information (E));
55
          end loop;
56
```

(continues on next page)

```
57 Close (F);
58 end Read_Exceptions;
59
60 end Exception_Occurrence_Stream;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Exceptions_Package.Exception_

Goccurrence_Stream

MD5: 3d9f2bd9480aa6dcc250b249b9ef4870
```

Runtime output

```
raised EXCEPTION_OCCURRENCE_STREAM.CUSTOM_EXCEPTION : We got a problem
```

raised CONSTRAINT_ERROR : Constraint is not correct

In this example, we store the exceptions raised in the application in the *exceptions_file.bin* file. In the exception part of procedures Nested_1 and Nested_2, we call Exception_Occurrence'Write to store an exception occurence in the file. In the Read_Exceptions block, we read the exceptions from the the file by calling Exception_Occurrence'Read.

12.5.3 Debugging exceptions in the GNAT toolchain

Here is a typical exception handler that catches all unexpected exceptions in the application:

Listing 44: main.adb

```
with Ada.Exceptions:
1
   with Ada.Text I0;
                         use Ada.Text I0;
2
3
   procedure Main is
4
5
      procedure Nested is
6
      begin
7
          raise Constraint Error
8
                  with "some message";
9
      end Nested:
10
11
   begin
12
      Nested;
13
14
   exception
15
      when E : others =>
16
          Put Line
17
           (Ada.Exceptions.Exception_Information (E));
18
   end Main:
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Exceptions_Package.Exception_

GINFormation

MD5: f95068ca90d79b92a7c2031322349153
```

Runtime output

raised CONSTRAINT_ERROR : some message

The output we get when running the application is not very informative. To get more information, we need to rerun the program in the debugger. To make the session more interesting though, we should add debug information in the executable, which means using the -g switch in the **gnatmake** command.

The session would look like the following (omitting some of the output from the debugger):

```
# Cleanup previous compilation
> rm *.o
> gnatmake -g main.adb
> gdb ./main
(gdb) catch exception
(gdb) run
Catchpoint 1, CONSTRAINT ERROR at 0x000000000402860 in main.nested () at main.
→adb:8
8
               raise Constraint Error with "some message";
(qdb) bt
#0 < gnat debug raise exception> (e=0x62ec40 <constraint error>) at s-excdeb.
→adb:43
#1 0x000000000040426f in ada.exceptions.complete occurrence (x=x@entry=0x637050)
at a-except.adb:934
#2 0x000000000040427b in ada.exceptions.complete and propagate occurrence (
x=x@entry=0x637050) at a-except.adb:943
#3 0x00000000004042d0 in <_ gnat_raise_exception> (e=0x62ec40 <constraint_error>,
message=...) at a-except.adb:982
#4 0x0000000000402860 in main.nested ()
#5 0x000000000040287c in main ()
```

And we now know exactly where the exception was raised. But in fact, we could have this information directly when running the application. For this, we need to bind the application with the switch -E, which tells the binder to store exception tracebacks in exception occurrences. Let's recompile and rerun the application.

```
> rm *.o # Cleanup previous compilation
> gnatmake -g main.adb -bargs -E
> ./main
Exception name: CONSTRAINT_ERROR
Message: some message
Call stack traceback locations:
0x10b7e24d1 0x10b7e24ee 0x10b7e2472
```

The traceback, as is, is not very useful. We now need to use another tool that is bundled with GNAT, called **addr2line**. Here is an example of its use:

```
> addr2line -e main --functions --demangle 0x10b7e24d1 0x10b7e24ee 0x10b7e2472
/path/main.adb:8
_ada_main
/path/main.adb:12
main
/path/b~main.adb:240
```

This time we do have a symbolic backtrace, which shows information similar to what we got in the debugger.

For users on OSX machines, **addr2line** does not exist. On these machines, however, an equivalent solution exists. You need to link your application with an additional switch, and then use the tool **atos**, as in:

```
> rm *.0
> gnatmake -g main.adb -bargs -E -largs -Wl,-no_pie
> ./main
Exception name: CONSTRAINT_ERROR
Message: some message
Call stack traceback locations:
0x1000014d1 0x1000014ee 0x100001472
> atos -o main 0x1000014d1 0x1000014ee 0x100001472
main__nested.2550 (in main) (main.adb:8)
_ada_main (in main) (main.adb:12)
main (in main) + 90
```

We will now discuss a relatively new switch of the compiler, namely -gnateE. When used, this switch will generate extra information in exception messages.

Let's amend our test program to:

Listing 45: main.adb

```
with Ada.Exceptions;
1
   with Ada.Text I0;
                            use Ada.Text I0;
2
3
   procedure Main is
4
5
      procedure Nested (Index : Integer) is
6
          type T_Array is array (1 .. 2) of Integer;
7
          T : constant T_Array := (10, 20);
8
      begin
9
         Put_Line (T (Index)'Img);
10
      end Nested;
11
12
   begin
13
      Nested (3);
14
15
   exception
16
      when E : others =>
17
          Put Line
18
           (Ada.Exceptions.Exception_Information (E));
19
   end Main;
20
```

Code block metadata

Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Exceptions_Package.Exception_ GINFormation MD5: 3590f2bf48f6ed1cf7745d576924cad4

Runtime output

```
raised CONSTRAINT_ERROR : main.adb:10:17 index check failed
index 3 not in 1..2
```

When running the application, we see that the exception information (traceback) is the same as before, but this time the exception message is set automatically by the compiler. So we know we got a Constraint_Error because an incorrect index was used at the named source location (main.adb, line 10). But the significant addition is the second line of the message, which indicates exactly the cause of the error. Here, we wanted to get the element at index 3, in an array whose range of valid indexes is from 1 to 2. (No need for a debugger in this case.)

The column information on the first line of the exception message is also very useful when dealing with null pointers. For instance, a line such as:

```
A := Rec1.Rec2.Rec3.Rec4.all;
```

where each of the Rec is itself a pointer, might raise Constraint_Error with a message "access check failed". This indicates for sure that one of the pointers is null, and by using the column information it is generally easy to find out which one it is.

12.6 Exception renaming

We can rename exceptions by using the an exception renaming declaration in this form Renamed_Exception : exception renames Existing_Exception;. For example:

Listing 46: show_exception_renaming.adb

```
procedure Show_Exception_Renaming is
CE : exception renames Constraint_Error;
hegin
```

```
3 begi
```

```
4 raise CE;
5 end Show_Exception_Renaming;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Exception_Renaming.Exception_

→Renaming

MD5: ff20825162ee9eef6ac8ed329da2a80f
```

Runtime output

```
raised CONSTRAINT_ERROR : show_exception_renaming.adb:4
```

Exception renaming creates a new view of the original exception. If we rename an exception from package A in package B, that exception will become visible in package B. For example:

Listing 47: internal_exceptions.ads

```
1 package Internal_Exceptions is
2
3 Int_E : exception;
4
5 end Internal_Exceptions;
```

Listing 48: test_constraints.ads

```
with Internal_Exceptions;
package Test_Constraints is
Ext_E : exception renames
Internal_Exceptions.Int_E;
end Test_Constraints;
```

Listing 49: show exception renaming view.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Ada.Exceptions; use Ada.Exceptions;
with Test_Constraints; use Test_Constraints;
```

(continues on next page)

```
procedure Show_Exception_Renaming_View is
6
   begin
7
      raise Ext E;
8
   exception
9
      when E : others =>
10
         Put Line
11
           (Ada.Exceptions.Exception_Information (E));
12
  end Show_Exception_Renaming_View;
13
```

Code block metadata

Runtime output

raised INTERNAL_EXCEPTIONS.INT_E : show_exception_renaming_view.adb:8

Here, we're renaming the Int_E exception in the Test_Constraints package. The Int_E exception isn't directly visible in the Show_Exception_Renaming procedure because we're not withing the Internal_Exceptions package. However, it is indirectly visible in that procedure via the renaming (Ext_E) in the Test_Constraints package.

```
In the Ada Reference Manual
```

• 8.5.2 Exception Renaming Declarations²³⁹

12.7 Out and Uninitialized

\rm 1 Note

This section was originally written by Robert Dewar and published as Gem #150: Out and Uninitialized²⁴⁰

Perhaps surprisingly, the Ada standard indicates cases where objects passed to **out** and **in out** parameters might not be updated when a procedure terminates due to an exception. Let's take an example:

Listing 50: show_out_uninitialized.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   procedure Show_Out_Uninitialized is
2
3
      procedure Local (A
                               : in out Integer;
4
                        Error : Boolean) is
5
      begin
6
         A := 1;
7
8
         if Error then
9
             raise Program_Error;
10
         end if;
11
```

(continues on next page)

²³⁹ http://www.ada-auth.org/standards/22rm/html/RM-8-5-2.html
 ²⁴⁰ https://www.adacore.com/gems/gem-150out-and-uninitialized

```
end Local;
12
13
      B : Integer := 0;
14
15
   begin
16
      Local (B, Error => True);
17
   exception
18
      when Program_Error =>
19
         Put_Line ("Value for B is"
20
                    & Integer'Image (B)); -- "0"
21
   end Show Out Uninitialized;
22
```

Code block metadata

```
Project: Courses.Advanced Ada.Control Flow.Exceptions.Out Uninitialized.Out
→Uninitialized 1
MD5: cebcf14e9fd088e38b98a5132d9fd998
```

Runtime output

Value for B is 0

This program outputs a value of 0 for B, whereas the code indicates that A is assigned before raising the exception, and so the reader might expect B to also be updated.

The catch, though, is that a compiler must by default pass objects of elementary types (scalars and access types) by copy and might choose to do so for other types (records, for example), including when passing out and in out parameters. So what happens is that while the formal parameter A is properly initialized, the exception is raised before the new value of A has been copied back into B (the copy will only happen on a normal return).

In the GNAT toolchain

In general, any code that reads the actual object passed to an **out** or **in out** parameter after an exception is suspect and should be avoided. GNAT has useful warnings here, so that if we simplify the above code to:

Listing 51: show out uninitialized warnings.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Out Uninitialized Warnings is
3
4
        procedure Local (A : in out Integer) is
5
        begin
6
           A := 1;
           raise Program Error;
8
        end Local;
9
10
      B : Integer := 0;
11
12
   begin
13
      Local (B);
14
   exception
15
      when others =>
16
          Put Line ("Value for B is"
17
                     & Integer'Image (B));
18
   end Show_Out_Uninitialized_Warnings;
19
```

Code block metadata

7

Build output

```
show_out_uninitialized_warnings.adb:7:10: warning: assignment to pass-by-copy_

oformal may have no effect [enabled by default]

show_out_uninitialized_warnings.adb:7:10: warning: "raise" statement may result_

of abnormal return (RM 6.4.1(17)) [enabled by default]
```

Runtime output

Value for B is 0

We now get a compilation warning that the pass-by-copy formal may have no effect.

Of course, GNAT is not able to point out all such errors (see first example above), which in general would require full flow analysis.

The behavior is different when using parameter types that the standard mandates be passed by reference, such as tagged types for instance. So the following code will work as expected, updating the actual parameter despite the exception:

Listing 52: show_out_initialized_rec.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Out Initialized Rec is
3
4
       type Rec is tagged record
5
         Field : Integer;
6
       end record;
7
8
       procedure Local (A : in out Rec) is
9
       begin
10
          A.Field := 1;
11
          raise Program Error;
12
       end Local:
13
14
       V : Rec:
15
16
   begin
17
       V.Field := 0;
18
       Local (V);
19
   exception
20
       when others =>
21
          Put Line ("Value of Field is"
22
                     & V.Field'Img); -- "1"
23
   end Show Out Initialized Rec;
24
```

Code block metadata

Runtime output

Value of Field is 1

```
In the GNAT toolchain
   It's worth mentioning that GNAT provides a pragma called Export Procedure that forces
   reference semantics on out parameters. Use of this pragma would ensure updates of the
   actual parameter prior to abnormal completion of the procedure. However, this pragma
   only applies to library-level procedures, so the examples above have to be rewritten to
   avoid the use of a nested procedure, and really this pragma is intended mainly for use
   in interfacing with foreign code. The code below shows an example that ensures that B
   is set to 1 after the call to Local:
                            Listing 53: exported procedures.ads
   package Exported Procedures is
1
2
     procedure Local (A
                            : in out Integer;
3
                       Error : Boolean);
4
     pragma Export Procedure
5
        (Local,
6
       Mechanism => (A => Reference));
7
8
   end Exported Procedures;
9
                            Listing 54: exported_procedures.adb
   package body Exported_Procedures is
1
2
      procedure Local (A
                             : in out Integer;
3
                        Error : Boolean) is
4
      begin A := 1;
5
          if Error then
6
            raise Program Error;
7
         end if;
8
      end Local;
9
10
   end Exported Procedures;
11
                             Listing 55: show out reference.adb
   with Ada.Text IO; use Ada.Text IO;
1
2
   with Exported Procedures;
3
   use Exported Procedures;
4
5
   procedure Show Out Reference is
6
      B : Integer := \overline{0};
7
   beain
8
      Local (B, Error => True);
9
   exception
10
      when Program Error =>
11
          Put_Line ("Value for B is"
12
                    & Integer'Image (B)); -- "1"
13
   end Show Out Reference;
14
   Code block metadata
   Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Out_Uninitialized.Out_

→Uninitialized 4
   MD5: aed2788be2b3ceeec19b28421c53fc66
   Runtime output
```

Value for B is 1

In the case of direct assignments to global variables, the behavior in the presence of exceptions is somewhat different. For predefined exceptions, most notably Constraint_Error, the optimization permissions allow some flexibility in whether a global variable is or is not updated when an exception occurs (see Ada RM 11.6²⁴¹). For instance, the following code makes an incorrect assumption:

X := 0; -- about to try addition Y := Y + 1; -- see if addition raises exception X := 1 -- addition succeeded

A program is not justified in assuming that X = 0 if the addition raises an exception (assuming X is a global here). So any such assumptions in a program are incorrect code which should be fixed.

In the Ada Reference Manual

• 11.6 Exceptions and Optimization²⁴²

12.8 Suppressing checks

12.8.1 pragma Suppress

1 Note

This section was originally written by Gary Dismukes and published as Gem #63: The Effect of Pragma Suppress²⁴³.

One of Ada's key strengths has always been its strong typing. The language imposes stringent checking of type and subtype properties to help prevent accidental violations of the type system that are a common source of program bugs in other less-strict languages such as C. This is done using a combination of compile-time restrictions (legality rules), that prohibit mixing values of different types, together with run-time checks to catch violations of various dynamic properties. Examples are checking values against subtype constraints and preventing dereferences of null access values.

At the same time, Ada does provide certain "loophole" features, such as Unchecked_Conversion, that allow selective bypassing of the normal safety features, which is sometimes necessary when interfacing with hardware or code written in other languages.

Ada also permits explicit suppression of the run-time checks that are there to ensure that various properties of objects are not violated. This suppression can be done using **pragma** *Suppress*, as well as by using a compile-time switch on most implementations — in the case of GNAT, with the -gnatp switch.

In addition to allowing all checks to be suppressed, **pragma** *Suppress* supports suppression of specific forms of check, such as Index_Check for array indexing, Range_Check for scalar bounds checking, and Access_Check for dereferencing of access values. (See section 11.5 of the Ada Reference Manual for further details.)

Here's a simple example of suppressing index checks within a specific subprogram:

²⁴¹ http://www.ada-auth.org/standards/22rm/html/RM-11-6.html

²⁴² http://www.ada-auth.org/standards/22rm/html/RM-11-6.html

²⁴³ https://www.adacore.com/gems/gem-63

```
procedure Main is
    procedure Sort_Array (A : in out Some_Array) is
    pragma Suppress (Index_Check);
    -- eliminate check overhead
    begin
    ...
    end Sort_Array;
end Main;
```

Unlike a feature such as Unchecked_Conversion, however, the purpose of check suppression is not to enable programs to subvert the type system, though many programmers seem to have that misconception.

What's important to understand about **pragma** *Suppress* is that it only gives permission to the implementation to remove checks, but doesn't require such elimination. The intention of Suppress is not to allow bypassing of Ada semantics, but rather to improve efficiency, and the Ada Reference Manual has a clear statement to that effect in the note in RM-11.5, paragraph 29:

There is no guarantee that a suppressed check is actually removed; hence a **pragma** *Suppress* should be used only for efficiency reasons.

There is associated Implementation Advice that recommends that implementations should minimize the code executed for checks that have been suppressed, but it's still the responsibility of the programmer to ensure that the correct functioning of the program doesn't depend on checks not being performed.

There are various reasons why a compiler might choose not to remove a check. On some hardware, certain checks may be essentially free, such as null pointer checks or arithmetic overflow, and it might be impractical or add extra cost to suppress the check. Another example where it wouldn't make sense to remove checks is for an operation implemented by a call to a run-time routine, where the check might be only a small part of a more expensive operation done out of line.

Furthermore, in many cases GNAT can determine at compile time that a given run-time check is guaranteed to be violated. In such situations, it gives a warning that an exception will be raised, and generates code specifically to raise the exception. Here's an example:

```
X : Integer range 1..10 := ...;
...
if A > B then
        X := X + 1;
...
end if;
```

For the assignment incrementing X, the compiler will normally generate machine code equivalent to:

```
Temp := X + 1;
if Temp > 10 then
    raise Constraint_Error;
end if;
X := Temp;
```

If range checks are suppressed, then the compiler can just generate the increment and assignment. However, if the compiler is able to somehow prove that X = 10 at this point, it will issue a warning, and replace the entire assignment with simply:

raise Constraint_Error;

even though checks are suppressed. This is appropriate, because

- 1. we don't care about the efficiency of buggy code, and
- 2. there is no "extra" cost to the check, because if we reach that point, the code will unconditionally fail.

One other important thing to note about checks and **pragma** *Suppress* is this statement in the Ada RM (RM-11.5, paragraph 26):

If a given check has been suppressed, and the corresponding error situation occurs, the execution of the program is erroneous.

In Ada, erroneous execution is a bad situation to be in, because it means that the execution of your program could have arbitrary nasty effects, such as unintended overwriting of memory. Note also that a program whose "correct" execution somehow depends on a given check being suppressed might work as the programmer expects, but could still fail when compiled with a different compiler, or for a different target, or even with a newer version of the same compiler. Other changes such as switching on optimization or making a change to a totally unrelated part of the code could also cause the code to start failing.

So it's definitely not wise to write code that relies on checks being removed. In fact, it really only makes sense to suppress checks once there's good reason to believe that the checks can't fail, as a result of testing or other analysis. Otherwise, you're removing an important safety feature of Ada that's intended to help catch bugs.

12.8.2 pragma Unsuppress

We can use **pragma** *Unsuppress* to reverse the effect of a **pragma** *Suppress*. While **pragma** *Suppress* gives permission to the compiler to remove a specific check, **pragma** *Unsuppress* revokes that permission.

Let's see an example:

```
Listing 56: show index check.adb
```

```
procedure Show_Index_Check is
1
2
      type Integer Array is
3
         array (Positive range <>) of Integer;
4
5
      pragma Suppress (Index_Check);
6
          from now on, the compiler may
7
           eliminate index checks...
8
       - -
9
      function Unchecked_Value_Of
10
         (A : Integer_Array;
11
          I : Integer)
12
          return Integer
13
      is
14
          type Half Integer Array is new
15
            Integer_Array (A'First ..
16
                            A'First + A'Length / 2);
17
18
19
         A_2 : Half_Integer_Array := (others => 0);
20
      begin
          return A_2 (I);
21
      end Unchecked_Value_Of;
22
23
      pragma Unsuppress (Index_Check);
24
```

(continues on next page)

```
-- from now on, index checks are
25
       -- typically performed...
26
27
       function Value_Of
28
         (A : Integer_Array;
29
          I : Integer)
30
          return Integer
31
       is
32
          type Half_Integer_Array is new
33
            Integer_Array (A'First ..
34
                             A'First + A'Length / 2);
35
36
          A_2 : Half_Integer_Array := (others => 0);
37
38
       begin
          return A_2 (I);
39
       end Value_0f;
40
41
       Arr_1 : Integer_Array (1 .. 10) :=
42
                  (others => 1);
43
44
   begin
45
       Arr_1 (10) := Unchecked_Value_Of (Arr_1, 10);
46
       Arr_1 (10) := Value_Of (Arr_1, 10);
47
48
   end Show_Index_Check;
49
```

Code block metadata

```
Project: Courses.Advanced_Ada.Control_Flow.Exceptions.Pragma_Unsuppress.Pragma_

→Unsuppress

MD5: 0585b78fd57913d3172c7ab1ea6f4864
```

Runtime output

```
raised CONSTRAINT_ERROR : show_index_check.adb:39 index check failed
```

In this example, we first use a **pragma** *Suppress* (Index_Check), so the compiler is allowed to remove the index check from the Unchecked_Value_Of function. (Therefore, depending on the compiler, the call to the Unchecked_Value_Of function may complete without raising an exception.) Of course, in this specific example, suppressing the index check masks a severe issue.

In contrast, an index check is performed in the Value_Of function because of the **pragma** *Unsuppress*. As a result, the index checks fails in the call to this function, which raises a Constraint_Error exception.

0	In	the	Ada	Reference	Manual
---	----	-----	-----	-----------	--------

11.5 Suppressing Checks²⁴⁴

²⁴⁴ http://www.ada-auth.org/standards/22rm/html/RM-11-5.html

Part III

Modular programming

CHAPTER THIRTEEN

PACKAGES

13.1 Package renaming

We've seen in the Introduction to Ada course that we can rename packages²⁴⁵.

```
In the Ada Reference Manual
```

• 10.1.1 Compilation Units - Library Units²⁴⁶

13.1.1 Grouping packages

A use-case that we haven't mentioned in that course is that we can apply package renaming to group individual packages into a common hierarchy. For example:

Listing 1: driver_m1.ads

```
package Driver_M1 is
    end Driver M1;
```

Listing 2: driver m2.ads

```
1 package Driver_M2 is
2
3 end Driver_M2;
```

Listing 3: drivers.ads

```
package Drivers
with Pure is
a
end Drivers;
```

Listing 4: drivers-m1.ads

```
with Driver_M1;
```

3

package Drivers.M1 renames Driver_M1;

²⁴⁵ https://learn.adacore.com/courses/intro-to-ada/chapters/modular_programming.html# intro-ada-package-renaming

246 http://www.ada-auth.org/standards/22rm/html/RM-10-1-1.html

Listing 5: drivers-m2.ads

```
with Driver_M2;
```

```
3
```

1

3 4

5

4

5

2

package Drivers.M2 renames Driver M2;

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Package_Renaming.Package_ →Renaming 1 MD5: 8d6a6bec32f7ec4397de1faf9f0b44d9

Here, we're renaming the Driver M1 and Driver M2 packages as child packages of the Drivers package, which is a pure package.

```
1 Important
   Note that a package that is renamed as a child package cannot refer to information from
   its (non-renamed) parent. In other words, Driver M1 (renamed as Drivers.M1) cannot
   refer to information from the Drivers package. For example:
                                 Listing 6: driver m1.ads
   package Driver M1 is
2
      Counter_2 : Integer := Drivers.Counter;
  end Driver M1;
                                  Listing 7: drivers.ads
   package Drivers is
1
2
      Counter : Integer := 0;
3
  end Drivers;
                                Listing 8: drivers-m1.ads
  with Driver M1;
1
  package Drivers.M1 renames Driver M1;
   Code block metadata
   Project: Courses.Advanced_Ada.Modular_Prog.Packages.Package_Renaming.Package_
    →Renaming_1_Refer_To_Parent
   MD5: d174746d8151d9a2cd048ad44e853850
   Build output
   driver m1.ads:3:27: error: "Drivers" is undefined
   qprbuild: *** compilation phase failed
```

As expected, compilation fails here because Drivers.Counter isn't visible in Driver M1, even though the renaming (Drivers.M1) creates a virtual hierarchy.

13.1.2 Child of renamed package

Note that we cannot create a child package using a parent package name that was introduced by a renaming. For example, let's say we want to create a child package Ext for the Drivers.M1 package we've seen earlier. We cannot just declare a Drivers.M1.Ext package like this:

```
package Drivers.M1.Ext is
```

```
end Drivers.M1.Ext;
```

because the parent unit cannot be a renaming. The solution is to actually extend the original (non-renamed) package:

Listing 9: driver_m1-ext.ads

```
package Driver_M1.Ext is
    end Driver_M1.Ext;
```

Listing 10: dummy.adb

```
-- A package called Drivers.M1.Ext is
1
       automatically available!
2
  - -
3
  with Drivers.M1.Ext:
4
5
  procedure Dummy is
6
  begin
7
      null;
8
  end Dummy;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Package_Renaming.Package_

⊲Renaming_1

MD5: e338d668dbd98b1a3917a8d3d948a439
```

This works fine because any child package of a package P is also a child package of a renamed version of P. (Therefore, because Ext is a child package of Driver_M1, it is also a child package of the renamed Drivers.M1 package.)

13.1.3 Backwards-compatibility via renaming

We can also use renaming to ensure backwards-compatibility when changing the package hierarchy. For example, we could adapt the previous source-code by:

- converting Driver_M1 and Driver_M2 to child packages of Drivers, and
- using package renaming to *mimic* the original names (Driver_M1 and Driver_M2).

This is the adapted code:

Listing 11: drivers.ads

```
package Drivers
with Pure is

end Drivers;
```

```
Listing 12: drivers-m1.ads
```

```
1 -- We've converted Driver_M1 to
2 -- Drivers.M1:
3
4 package Drivers.M1 is
5
6 end Drivers.M1;
```

Listing 13: drivers-m2.ads

```
1 -- We've converted Driver_M2 to
2 -- Drivers.M2:
3
4 package Drivers.M2 is
5
6 end Drivers.M2;
```

Listing 14: driver_m1.ads

```
1 -- Original Driver_M1 package still
2 -- available via package renaming:
3
4 with Drivers.M1;
5
6 package Driver M1 renames Drivers.M1;
```

Listing 15: driver_m2.ads

```
1 -- Original Driver_M2 package still
2 -- available via package renaming:
3
4 with Drivers.M2;
5
6 package Driver_M2 renames Drivers.M2;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Package_Renaming.Package_

⊸Renaming_2

MD5: 27f8066b5f5954514fea51b6e9b9de81
```

Now, M1 and M2 are *actual* child packages of Drivers, but their original names are still available. By doing so, we ensure that existing software that makes use of the original packages doesn't break.

13.2 Private packages

In this section, we discuss the concept of private packages. However, before we proceed with the discussion, let's recapitulate some important ideas that we've seen earlier.

In the Introduction to Ada course²⁴⁷, we've seen that encapsulation plays an important role in modular programming. By using the private part of a package specification, we can disclose some information, but, at the same time, prevent that this information gets accessed where it shouldn't be used directly. Similarly, we've seen that we can use the private part of a package to distinguish between the *partial and full view* (page 38) of a data type.

²⁴⁷ https://learn.adacore.com/courses/intro-to-ada/chapters/privacy.html#intro-ada-course-privacy

The main application of private packages is to create private child packages, whose purpose is to serve as internal implementation packages within a package hierarchy. By doing so, we can expose the internals to other public child packages, but prevent that external clients can directly access them.

As we'll see next, there are many rules that ensure that internal visibility is enforced for those private child packages. At the same time, the same rules ensure that private packages aren't visible outside of the package hierarchy.

13.2.1 Declaration and usage

We declare private packages by using the **private** keyword. For example, let's say we have a package named Data_Processing:

```
Listing 16: data processing.ads
```

```
1 package Data_Processing is
2
3 -- ...
4
5 end Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

→Package_Decl

MD5: 502811212890785d90c6f891d7f8e557
```

We simply write **private package** to declare a private child package named Calculations:

Listing 17: data_processing-calculations.ads

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

⊲Package_Decl

MD5: 20df8b2ac4c9aa93f03a12afd9b7ef30
```

Let's see a complete example:

```
Listing 18: data_processing.ads
```

```
package Data_Processing is
1
2
      type Data is private;
3
4
      procedure Process (D : in out Data);
5
6
   private
7
8
9
       type Data is null record;
10
   end Data_Processing;
11
```

Listing 19: data_processing-calculations.ads

```
private package Data_Processing.Calculations is
procedure Calculate (D : in out Data);
end Data_Processing.Calculations;
```

Listing 20: data_processing.adb

```
with Data Processing.Calculations;
1
   use Data_Processing.Calculations;
2
3
   package body Data_Processing is
4
5
      procedure Process (D : in out Data) is
6
      begin
7
         Calculate (D);
8
      end Process;
9
10
   end Data_Processing;
11
```

Listing 21: data_processing-calculations.adb

```
package body Data_Processing.Calculations is
1
2
      procedure Calculate (D : in out Data) is
3
      begin
4
             Dummy implementation...
5
6
         null;
      end Calculate;
7
8
  end Data_Processing.Calculations;
9
```

Listing 22: test_data_processing.adb

```
with Data_Processing; use Data_Processing;
procedure Test_Data_Processing is
    D : Data;
begin
Process (D);
end Test_Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

→Package

MD5: 3edd5f73938e809994347b5876014d0d
```

In this example, we refer to the private child package Calculations in the body of the Data_Processing package — by simply writing with Data_Processing.Calculations. After that, we can call the Calculate procedure normally in the Process procedure.

13.2.2 Private sibling packages

We can introduce another private package Advanced_Calculations as a child of Data_Processing and refer to the Calculations package in its specification:

Listing 23: data_processing.ads

```
package Data_Processing is
1
2
      type Data is private;
3
4
      procedure Process (D : in out Data);
5
6
   private
7
8
      type Data is null record;
9
10
   end Data_Processing;
11
```

Listing 24: data_processing-calculations.ads

```
private package Data_Processing.Calculations is
procedure Calculate (D : in out Data);
end Data_Processing.Calculations;
```

Listing 25: data_processing-advanced_calculations.ads

```
with Data_Processing.Calculations;
1
   use Data_Processing.Calculations;
2
3
   private
4
   package Data Processing.Advanced Calculations is
5
6
      procedure Advanced Calculate (D : in out Data)
7
        renames Calculate;
8
9
   end Data Processing.Advanced Calculations;
10
```

Listing 26: data_processing.adb

```
with Data_Processing.Advanced_Calculations;
1
   use Data_Processing.Advanced_Calculations;
2
3
   package body Data_Processing is
4
5
      procedure Process (D : in out Data) is
6
      begin
7
         Advanced_Calculate (D);
8
      end Process;
9
10
   end Data_Processing;
11
```

Listing 27: data_processing-calculations.adb

```
package body Data_Processing.Calculations is
1
2
      procedure Calculate (D : in out Data) is
3
      begin
4
             Dummy implementation...
5
         null;
6
      end Calculate;
7
8
  end Data_Processing.Calculations;
9
```

Listing 28: test_data_processing.adb

```
with Data_Processing; use Data_Processing;
procedure Test_Data_Processing is
D : Data;
begin
Process (D);
end Test_Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

→Package_2

MD5: 32fc76ae13f1eecdd854a029793034d8
```

Note that, in the body of the Data_Processing package, we're now referring to the new Advanced_Calculations package instead of the Calculations package.

Referring to a private child package in the specification of another private child package is OK, but we cannot do the same in the specification of a *non-private* package. For example, let's change the specification of the Advanced Calculations and make it *non-private*:

Listing 29: data_processing-advanced_calculations.ads

```
with Data_Processing.Calculations;
use Data_Processing.Calculations;
package Data_Processing.Advanced_Calculations is
procedure Advanced_Calculate (D : in out Data)
renames Calculate;
end Data Processing.Advanced Calculations;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

→Package_2

MD5: 27fd3bdb063a11ed7797cc44fa1e8349
```

Build output

```
data_processing-advanced_calculations.ads:1:06: error: current unit must also be

□ oprivate descendant of "Data_Processing"

gprbuild: *** compilation phase failed
```

Now, the compilation doesn't work anymore. However, we could still refer to Calculations packages in the body of the Advanced_Calculations package:

Listing 30: data_processing-advanced_calculations.ads

```
package Data_Processing.Advanced_Calculations is
procedure Advanced_Calculate (D : in out Data);
end Data_Processing.Advanced_Calculations;
```

Listing 31: data_processing-advanced_calculations.adb

```
with Data Processing.Calculations;
1
   use Data_Processing.Calculations;
2
3
   package body Data Processing. Advanced Calculations
4
   is
5
6
      procedure Advanced_Calculate (D : in out Data)
7
      is
8
      begin
9
        Calculate (D);
10
      end Advanced_Calculate;
11
12
   end Data Processing.Advanced Calculations;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

⊶Package_2

MD5: 3f37c129a6994c6b71a25ad17dcb440e
```

This works fine as expected: we can refer to private child packages in the body of another package — as long as both packages belong to the same package tree.

13.2.3 Outside the package tree

While we can use a with-clause of a private child package in the body of the Data_Processing package, we cannot do the same outside the package tree. For example, we cannot refer to it in the Test_Data_Processing procedure:

Listing 32: test_data_processing.adb

```
with Data_Processing; use Data_Processing;
1
2
  with Data Processing.Calculations;
3
   use Data_Processing.Calculations;
4
5
  procedure Test_Data_Processing is
6
      D : Data;
7
  begin
8
9
      Calculate (D);
  end Test_Data_Processing;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

→Package

MD5: c844327995b28d60c9a79b138a0f21d2
```

Build output

As expected, we get a compilation error because Calculations is only accessible within the Data_Processing, but not in the Test_Data_Processing procedure.

The same restrictions apply to child packages of private packages. For example, if we

implement a child package of the Calculations package — let's name it Calculations. Child —, we cannot refer to it in the Test_Data_Processing procedure:

```
Listing 33: data processing-calculations-child.ads
```

```
package Data_Processing.Calculations.Child is
procedure Process (D : in out Data);
end Data_Processing.Calculations.Child;
```

Listing 34: data_processing-calculations-child.adb

```
package body Data_Processing.Calculations.Child is
procedure Process (D : in out Data) is
begin
Calculate (D);
end Process;
end Data_Processing.Calculations.Child;
```

Listing 35: test_data_processing.adb

```
with Data_Processing; use Data_Processing;
1
2
   with Data Processing.Calculations.Child;
3
   use Data Processing.Calculations.Child;
4
5
  procedure Test Data Processing is
6
      D : Data;
7
8
  begin
      Calculate (D);
9
  end Test_Data_Processing;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_Packages.Private_

→Package

MD5: 2eaf23ddbab72578246ac07424008d9d
```

Build output

Again, as expected, we get an error because Calculations.Child — being a child of a private package — has the same restricted view as its parent package. Therefore, it cannot be visible in the Test_Data_Processing procedure as well. We'll discuss more about visibility *later* (page 568).

Note that subprograms can also be declared private. We'll see this *in another section* (page 585).

1 Important

We've discussed package renaming in a previous section (page 549). We can rename a package as a private package, too. For example: Listing 36: driver m1.ads package Driver M1 is 1 2 end Driver_M1; 3 Listing 37: drivers.ads package Drivers 1 with Pure is 2 3 4 end Drivers; Listing 38: drivers-m1.ads with Driver M1; 1 2 private package Drivers.M1 renames Driver M1; 3 Code block metadata Project: Courses.Advanced Ada.Modular Prog.Packages.Private Packages.Private →Package Renaming MD5: c03584dc26abb108c9c04074234b9637 Obviously, Drivers.M1 has the same restrictions as any private package: Listing 39: test driver.adb with Driver M1; 1 with Drivers.M1; 2 3 procedure Test Driver is 4 begin 5 null; 6 end Test_Driver; Code block metadata Project: Courses.Advanced Ada.Modular_Prog.Packages.Private_Packages.Private_ →Package Renaming MD5: 55415978604ccea4eeaeb02df13cd2f4 **Build output** test driver.adb:2:06: error: unit in with clause is private child unit test driver.adb:2:06: error: current unit must also have parent "Drivers" gprbuild: *** compilation phase failed As expected, although we can have the Driver M1 package in a with clause of the Test Driver procedure, we cannot do the same in the case of the Drivers.M1 package because it is private.

In the Ada Reference Manual

• 10.1.1 Compilation Units - Library Units²⁴⁸

²⁴⁸ http://www.ada-auth.org/standards/22rm/html/RM-10-1-1.html

13.3 Private with clauses

13.3.1 Definition and usage

A private with clause allows us to refer to a package in the private part of another package. For example, if we want to refer to package P in the private part of Data, we can write **private with** P:

Listing 40: p.ads

```
1 package P is
2
3 type T is null record;
4
5 end P;
```

Listing 41: data.ads

```
private with P;
1
2
   package Data is
3
4
      type T2 is private;
5
6
7
   private
8
       -- Information from P is
9
       -- visible here
10
      type T2 is new P.T;
11
12
   end Data;
13
```

Listing 42: main.adb

```
with Data; use Data;
procedure Main is
A : T2;
begin
null;
r end Main;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_With_Clauses.Simple_

⊶Private_With_Clause

MD5: d0705add0dd7861c83822b0d35dacba4
```

As you can see in the example, as the information from P is available in the private part of Data, we can derive a new type T2 based on T from P. However, we cannot do the same in the visible part of Data:

Listing 43: data.ads

```
private with P;
package Data is
-- ERROR: information from P
-- isn't visible here
```

(continues on next page)

8 type T2 is new P.T;

```
9
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_With_Clauses.Simple_

⇔Private_With_Clause

MD5: b454e875f73432f5632a20ab40ae7da6
```

Build output

end Data;

```
data.ads:8:19: error: "P" is not visible
data.ads:8:19: error: non-visible declaration at p.ads:1
gprbuild: *** compilation phase failed
```

Also, the information from P is available in the package body. For example, let's declare a Process procedure in the P package and use it in the body of the Data package:

Listing 44: p.ads

```
package P is
type T is null record;
procedure Process (A : T) is null;
end P;
```

Listing 45: data.ads

```
private with P;
1
2
   package Data is
3
4
      type T2 is private;
5
6
      procedure Process (A : T2);
7
8
   private
9
10
       -- Information from P is
11
       -- visible here
12
      type T2 is new P.T;
13
14
   end Data;
15
```

Listing 46: data.adb

```
1 package body Data is
2
3 procedure Process (A : T2) is
4 begin
5 P.Process (P.T (A));
6 end Process;
7
8 end Data;
```

Listing 47: main.adb

```
with Data; use Data;
procedure Main is
    A : T2;
begin
    null;
    end Main;
```

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_With_Clauses.Simple_ ⊶Private_With_Clause MD5: cecc09f95bd43dd7fd34a9e289bd2674

In the body of the Data, we can access information from the P package — as we do in the P.Process (P.T (A)) statement of the Process procedure.

13.3.2 Referring to private child package

There's one case where using a private with clause is the only way to refer to a package: when we want to refer to a private child package in another child package. For example, here we have a package P and its two child packages: **Private**_Child and Public_Child:

Listing 48: p.ads

Listing 49: p-private child.ads

```
1 private package P.Private_Child is
2
3 type T is null record;
4
5 end P.Private Child;
```

Listing 50: p-public_child.ads

```
private with P.Private Child;
1
2
   package P.Public_Child is
3
4
      type T2 is private;
5
6
   private
7
8
      type T2 is new P.Private_Child.T;
9
10
   end P.Public Child;
11
```

Listing 51: test_parent_child.adb

```
with P.Public_Child; use P.Public_Child;
procedure Test_Parent_Child is
A : T2;
```

(continues on next page)

```
5 begin
```

```
6 null;
```

```
7 end Test_Parent_Child;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Private_With_Clauses.Private_

⊲With_Clause

MD5: a6028416a957184be55a54f96a319e61
```

In this example, we're referring to the P.**Private**_Child package in the P.Public_Child package. As expected, this works fine. However, using a *normal* with clause doesn't work in this case:

Listing 52: p-public_child.ads

```
with P.Private Child;
1
2
   package P.Public Child is
3
4
      type T2 is private;
5
6
   private
7
8
       type T2 is new P.Private_Child.T;
9
10
   end P.Public Child;
11
```

Code block metadata

Build output

```
p-public_child.ads:1:06: error: current unit must also be private descendant of "P"
gprbuild: *** compilation phase failed
```

This gives an error because the information from the P.**Private**_Child, being a private child package, cannot be accessed in the public part of another child package. In summary, unless both packages are private packages, it's only possible to access the information from a private package in the private part of a non-private child package.

In the Ada Reference Manual

• 10.1.2 Context Clauses - With Clauses²⁴⁹

13.4 Limited Visibility

Sometimes, we might face the situation where two packages depend on information from each other. Let's consider a package A that depends on a package B, and vice-versa:

²⁴⁹ http://www.ada-auth.org/standards/22rm/html/RM-10-1-2.html

```
Listing 53: a.ads
```

```
with B; use B;
1
2
   package A is
3
4
      type T1 is record
5
         Value : T2;
6
      end record;
7
8
  end A;
9
```

Listing 54: b.ads

```
with A; use A;
1
2
   package B is
3
4
      type T2 is record
5
         Value : T1;
6
      end record;
7
8
9
```

end B;

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Limited_Visibility.Circular_
→Dependency
MD5: ae64f33706f1c58603aff2c33b02c910
```

Build output

```
a.ads:1:06: error: circular unit dependency
a.ads:1:06: error: "A (spec)" depends on "B (spec)"
a.ads:1:06: error: "B (spec)" depends on "A (spec)"
a.ads:1:06: error: "A (spec)" depends on "A (spec)"
gprbuild: *** compilation phase failed
```

Here, we have two *mutually dependent types* (page 177) T1 and T2, which are declared in two packages A and B that refer to each other. These with clauses constitute a circular dependency, so the compiler cannot compile either of those packages.

One way to solve this problem is by transforming this circular dependency into a partial dependency. We do this by limiting the visibility — using a limited with clause. To use a limited with clause for a package P, we simply write **limited with** P.

If a package A has limited visibility to a package B, then all types from package B are visible as if they had been declared as *incomplete types* (page 36). For the specific case of the previous source-code example, this would be the limited visibility to package B from package A's perspective:

```
package B is
   -- Incomplete type
  type T2;
```

end B;

As we've seen previously,

• we cannot declare objects of incomplete types, but we can declare access types and anonymous access objects of incomplete types. Also,

• we can use anonymous access types to declare *mutually dependent types* (page 177).

Keeping this information in mind, we can now correct the previous code by using limited with clauses for package A and declaring the component of the T1 record using an anonymous access type:

```
Listing 55: a.ads
```

```
1 limited with B;
2
3 package A is
4
5 type T1 is record
6 Ref : access B.T2;
7 end record;
8
9 end A;
```

Listing 56: b.ads

```
with A; use A;
package B is
type T2 is record
Value : T1;
end record;
end B;
```

```
Code block metadata
```

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Limited_Visibility.Limited_

⇔Visibility

MD5: 48591850665085a6fbb184f51b658a1b
```

As expected, we can now compile the code without issues.

Note that we can also use limited with clauses for both packages. If we do that, we must declare all components using anonymous access types:

Listing 57: a.ads

```
1 limited with B;
2
3 package A is
4
5 type T1 is record
6 Ref : access B.T2;
7 end record;
8
9 end A;
```

Listing 58: b.ads

```
1 limited with A;
2
3 package B is
4
5 type T2 is record
6 Ref : access A.T1;
7 end record;
```

8

```
end B;
```

Code block metadata

Now, both packages A and B have limited visibility to each other.

1 In the Ada Reference Manual

• 10.1.2 Context Clauses - With Clauses²⁵⁰

13.4.1 Limited visibility and private with clauses

We can limit the visibility and use *private with clauses* (page 560) at the same time. For a package P, we do this by simply writing **limited private with** P.

Let's reuse the previous source-code example and convert types T1 and T2 to private types:

Listing 59: a.ads

```
limited private with B;
1
2
   package A is
3
4
      type T1 is private;
5
6
   private
7
8
       -- Here, we have limited visibility
9
       -- of package B
10
11
      type T1 is record
12
         Ref : access B.T2;
13
      end record;
14
15
   end A;
16
```

Listing 60: b.ads

```
private with A;
1
2
   package B is
3
4
      type T2 is private;
5
6
   private
7
8
      use A;
9
10
       -- Here, we have full visibility
11
       -- of package A
12
13
      type T2 is record
14
```

(continues on next page)

²⁵⁰ http://www.ada-auth.org/standards/22rm/html/RM-10-1-2.html

```
15 Value : T1;
16 end record;
17
18 end B;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Limited_Visibility.Limited_

⊶Private_Visibility

MD5: b3ac546e2f55fb91229e834ca7a9783d
```

In this updated version of the source-code example, we have not only limited visibility to package B, but also, each package is just visible in the private part of the other package.

13.4.2 Limited visibility and other elements

It's important to mention that the limited visibility we've been discussing so far is restricted to type declarations — which are seen as incomplete types. In fact, when we use a limited with clause, all other declarations have no visibility at all! For example, let's say we have a package Info that declares a constant Zero_Const and a function Zero_Func:

Listing 61: info.ads

```
package Info is
function Zero_Func return Integer is (0);
Zero_Const : constant := 0;
end Info;
```

Code block metadata

Also, let's say we want to use the information (from package Info) in package A. If we have limited visibility to package Info, however, this information won't be visible. For example:

Listing 62: a.ads

```
limited private with Info;
1
2
   package A is
3
4
       type T1 is private;
5
6
   private
7
8
       type T1 is record
9
          V : Integer := Info.Zero_Const;
10
          W : Integer := Info.Zero Func;
11
      end record:
12
13
   end A;
14
```

Code block metadata

Build output

a.ads:10:26: error: "Zero_Const" not declared in "Info" a.ads:11:26: error: "Zero_Func" not declared in "Info" gprbuild: *** compilation phase failed

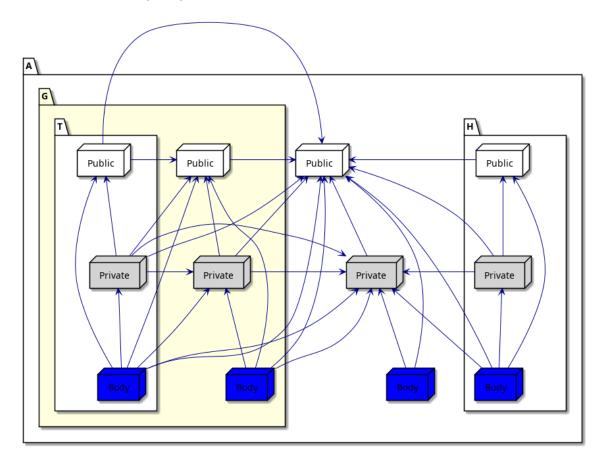
As expected, compilation fails because of the limited visibility — as Zero_Const and Zero_Func from the Info package are not visible in the private part of A. (Of course, if we revert to full visibility by simply removing the **limited** keyword from the example, the code compiles just fine.)

13.5 Visibility

In the previous sections, we already discussed visibility from various angles. However, it can be interesting to recapitulate this information with the help of diagrams that illustrate the different parts of a package and its relation with other units.

13.5.1 Automatic visibility

First, let's consider we have a package A, its children (A.G and A.H), and the grandchild A.G.T. As we've seen before, information of a parent package is automatically visible in its children. The following diagrams illustrates this:



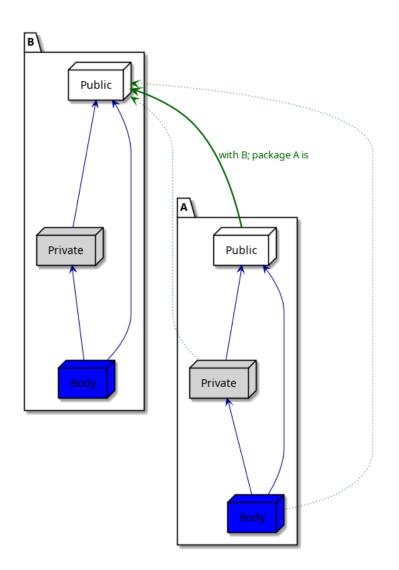
Because of this automatic visibility, many with clauses would be redundant in child pack-

ages. For example, we don't have to write with A; package A.G is, since the specification of package A is already visible in its child packages.

If we focus on package A.G (highlighted in the figure above), we see that it only has automatic visibility to its parent A, but not its child A.G.T. Also, it doesn't have visibility to its sibling A.H.

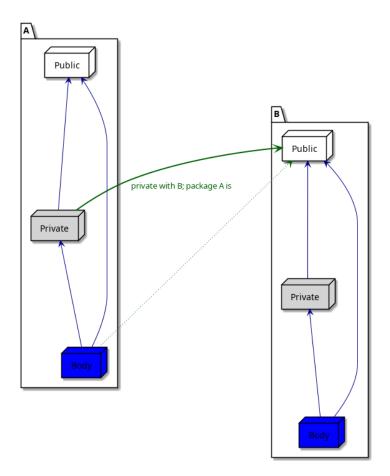
13.5.2 With clauses and visibility

In the rest of this section, we discuss all the situations where using with clauses is necessary to access the information of a package. Let's consider this example where we refer to a package B in the specification of a package A (using with B):

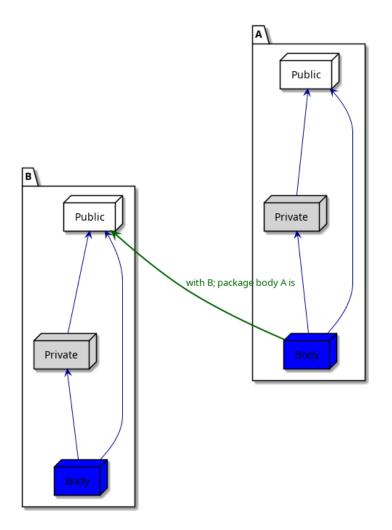


As we already know, the information from the public part of package B is visible in the public part of package A. In addition to that, it's also visible in the private part and in the body of package A. This is indicated by the dotted green arrows in the figure above.

Now, let's see the case where we refer to package B in the private part of package A (using **private with** B):



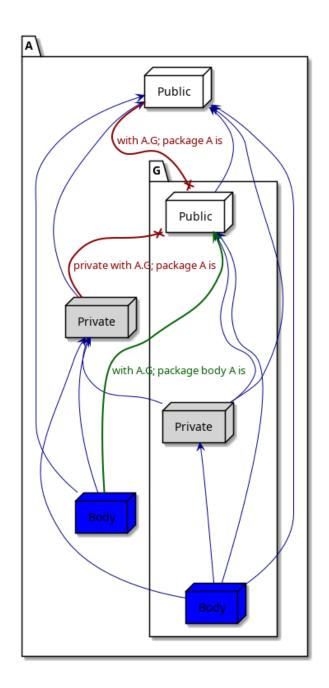
Here, the information is visible in the private part of package A, as well as in its body. Finally, let's see the case where we refer to package B in the body of package A:



Here, the information is only visible in the body of package A.

13.5.3 Circular dependency

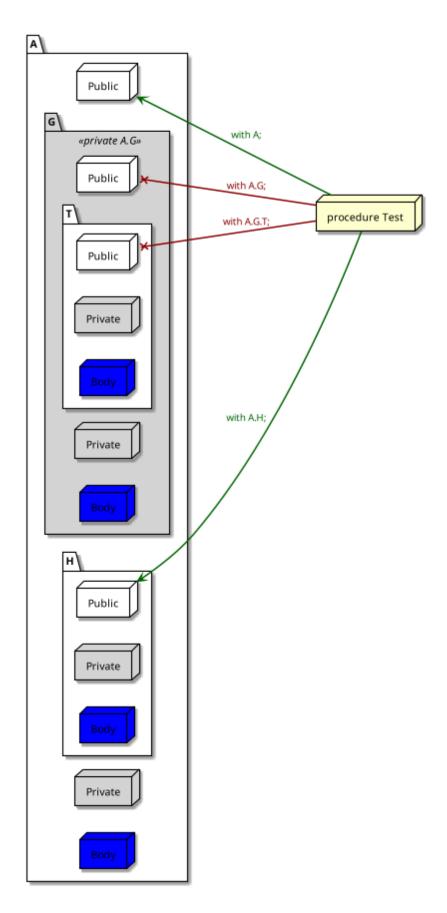
Let's return to package A and its descendants. As we've seen in previous sections, we cannot refer to a child package in the specification of its parent package because that would constitute circular dependency. (For example, we cannot write **with** A.G; **package** A is.) This situation — which causes a compilation error — is indicated by the red arrows in the figure below:



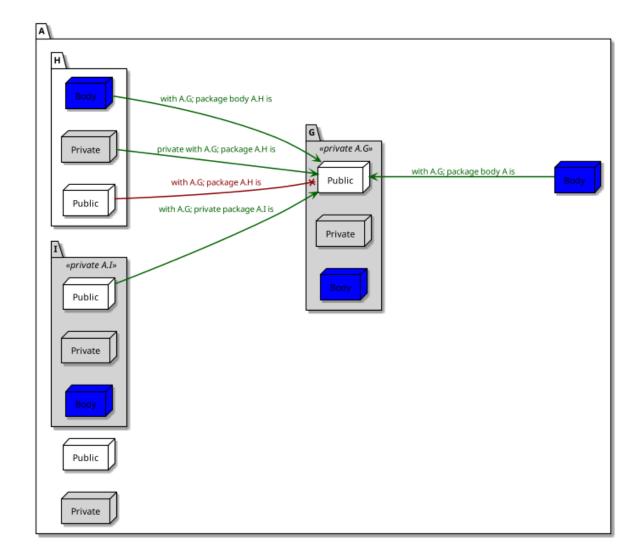
Note that referring to the child package A.G in the body of its parent is perfectly fine.

13.5.4 Private packages

The previous examples of this section only showed public packages. As we've seen before, we cannot refer to private packages outside of a package hierarchy, as we can see in the following example where we try to refer to package A and its descendants in the Test procedure:



As indicated by the red arrows, we cannot refer to the private child packages of A in the Test procedure, only the public child packages. Within the package hierarchy itself, we



cannot refer to the private package A.G in public sibling packages. For example:

Here, we cannot refer to the private package A.G in the public package A.H — as indicated by the red arrow. However, we can refer to the private package A.G in other private packages, such as A.I — as indicated by the green arrows.

13.6 Use type clause

Back in the Introduction to Ada course²⁵¹, we saw that use clauses provide direct visibility - in the scope where they're used - to the content of a package's visible part.

For example, consider this simple procedure:

Listing 63: display message.adb

```
with Ada.Text_I0;
procedure Display_Message is
begin
Ada.Text_I0.Put_Line ("Hello World!");
end Display_Message;
```

²⁵¹ https://learn.adacore.com/courses/intro-to-ada/chapters/modular_programming.html#intro-ada-use-clause

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.No_Use_Clause
MD5: 4c6ff19809c13ebd2fdfda482914e5f8

Runtime output

```
Hello World!
```

By adding **use** Ada.Text_I0 to this code, we make the visible part of the Ada.Text_I0 package directly visible in the scope of the Display_Message procedure, so we can now just write Put_Line instead of Ada.Text_I0.Put_Line:

Listing 64: display message.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Display_Message is
begin
Put_Line ("Hello World!");
end Display_Message;
```

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.Use_Clause
MD5: b105a777a1afd79008f8580cda432cfe

Runtime output

Hello World!

In this section, we discuss another example of use clauses. In addition, we introduce two specific forms of use clauses: **use** type and **use** all **type**.

In the Ada Reference Manual

```
• 8.4 Use Clauses<sup>252</sup>
```

13.6.1 Another use clause example

Let's now consider a simple package called Points, which contains the declaration of the Point type and two primitive: an Init function and an addition operator.

Listing 65: points.ads

```
package Points is
1
2
      type Point is private;
3
4
       function Init return Point;
5
6
       function "+" (P : Point;
7
                      I : Integer) return Point;
8
9
   private
10
11
       type Point is record
12
          X, Y : Integer;
13
```

(continues on next page)

²⁵² http://www.ada-auth.org/standards/22rm/html/RM-8-4.html

```
end record;
function Init return Point is (0, 0);
function "+" (P : Point;
I : Integer) return Point is
(P.X + I, P.Y + I);
end Points;
```

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.Use_Type_Clause
MD5: 1a43740d7231a3cc497e778866a12c55

We can implement a simple procedure that makes use of this package:

Listing 66: show point.adb

```
with Points; use Points;

procedure Show_Point is
    P : Point;

begin
    P := Init;
    P := P + 1;
end Show_Point;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.Use_Type_Clause
MD5: f5d44dd1fee8cf4d1a7e730f9a7c64cc
```

Here, we have a use clause, so we have direct visibility to the content of Points's visible part.

13.6.2 Visibility and Readability

In certain situations, however, we might want to avoid the use clause. If that's the case, we can rewrite the previous implementation by removing the use clause and specifying the Points package in the prefixed form:

Listing 67: show point.adb

```
with Points;
procedure Show_Point is
P : Points.Point;
begin
P := Points.Init;
P := Points."+" (P, 1);
end Show_Point;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.Use_Type_Clause
MD5: ca896b456a90c19b29ec4f262144c131
```

Although this code is correct, it might be difficult to read, as we have to specify the package whenever we're referring to a type or a subprogram from that package. Even worse: we now have to write operators in the prefixed form — such as Points."+" (P, 1).

13.6.3 use type

As a compromise, we can have direct visibility to the operators of a certain type. We do this by using a use clause in the form **use** type. This allows us to simplify the previous example:

Listing 68: show_point.adb

```
with Points:
1
2
   procedure Show Point is
3
      use type Points.Point;
4
5
      P : Points.Point;
6
   begin
7
      P := Points.Init;
8
      P := P + 1;
9
   end Show_Point;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.Use_Type_Clause
MD5: a9527276c27a67be8b5a59efcf6e5cfd
```

Note that **use** type just gives us direct visibility to the operators of a certain type, but not other primitives. For this reason, we still have to write Points.Init in the code example.

13.6.4 use all type

If we want to have direct visibility to all primitives of a certain type (and not just its operators), we need to write a use clause in the form **use** all **type**. This allows us to simplify the previous example even further:

Listing 69: show_point.adb

```
with Points:
1
2
   procedure Show Point is
3
      use all type Points.Point;
4
5
      P : Points.Point;
6
   begin
7
      P := Init;
8
      P := P + 1;
9
  end Show_Point;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Type_Clause.Use_Type_Clause
MD5: 4a8f6edd4e1811c4e8acb24393690282
```

Now, we've removed the prefix from all operations on the P variable.

13.7 Use clauses and naming conflicts

Visibility issues may arise when we have multiple use clauses. For instance, we might have types with the same name declared in multiple packages. This constitutes a naming conflict; in this case, the types become hidden — so they're not directly visible anymore, even if we have a use clause.

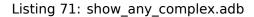


13.7.1 Code example

Let's start with a code example. First, we declare and implement a generic procedure that shows the value of a Complex object:

Listing 70: show_any_complex.ads

```
with Ada.Numerics.Generic_Complex_Types;
generic
with package Complex_Types is new
Ada.Numerics.Generic_Complex_Types (<>);
procedure Show_Any_Complex
(Msg : String;
Val : Complex_Types.Complex);
```



```
with Ada.Text I0;
1
   with Ada.Text_IO.Complex_IO;
2
3
   procedure Show Any Complex
4
      (Msg : String;
5
      Val : Complex Types.Complex)
6
   is
7
8
      package Complex_Float_Types_I0 is new
        Ada.Text_IO.Complex_IO (Complex_Types);
9
      use Complex_Float_Types_I0;
10
11
      use Ada.Text I0;
12
   begin
13
      Put (Msg & " ");
14
      Put (Val);
15
      New Line;
16
   end Show_Any_Complex;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Clause_Naming_Conflicts. GUse_Type_Clause_Complex_Types MD5: 2527291906d3a600eecd6d36e4359c1a

Then, we implement a test procedure where we declare the Complex_Float_Types package as an instance of the **Generic**_Complex_Types package:

Listing 72: show use.adb

```
with Ada.Numerics; use Ada.Numerics;
with Ada.Numerics.Generic_Complex_Types;
with Show_Any_Complex;
procedure Show_Use is
```

²⁵³ http://www.ada-auth.org/standards/22rm/html/RM-8-4.html

```
package Complex_Float_Types is new
8
         Ada.Numerics.Generic_Complex_Types
9
           (Real => Float);
10
      use Complex_Float_Types;
11
12
      procedure Show_Complex_Float is new
13
         Show_Any_Complex (Complex_Float_Types);
14
15
      C, D, X : Complex;
16
   begin
17
      C := Compose_From_Polar (3.0, Pi / 2.0);
18
      D := Compose_From_Polar (5.0, Pi / 2.0);
19
20
      X := C + D;
21
      Show_Complex_Float ("C:", C);
22
      Show_Complex_Float ("D:", D);
23
      Show_Complex_Float ("X:", X);
24
   end Show_Use;
25
```

Code block metadata

Runtime output

C: (-1.31134E-07, 3.00000E+00) D: (-2.18557E-07, 5.00000E+00) X: (-3.49691E-07, 8.00000E+00)

In this example, we declare variables of the Complex type, initialize them and use them in operations. Note that we have direct visibility to the package instance because we've added a simple use clause after the package instantiation — see **use** Complex_Float_Types in the example.

13.7.2 Naming conflict

Now, let's add the declaration of the Complex_Long_Float_Types package — a second instantiation of the **Generic**_Complex_Types package — to the code example:

Listing 73: show_use.adb

```
with Ada.Numerics; use Ada.Numerics;
1
2
   with Ada.Numerics.Generic_Complex_Types;
3
4
   with Show_Any_Complex;
5
6
   procedure Show Use is
7
      package Complex Float Types is new
8
         Ada.Numerics.Generic Complex Types
9
           (Real => Float);
10
      use Complex Float Types;
11
12
      package Complex Long Float Types is new
13
         Ada.Numerics.Generic_Complex_Types
14
           (Real => Long_Float);
15
      use Complex_Long_Float_Types;
16
17
```

```
procedure Show_Complex_Float is new
18
          Show_Any_Complex (Complex_Float_Types);
19
20
       C, D, X : Complex;
21
                   ^ ERROR: type is hidden!
22
       - -
23
   begin
       C := Compose_From_Polar (3.0, Pi / 2.0);
24
       D := Compose_From_Polar (5.0, Pi / 2.0);
25
       X := C + D;
26
27
       Show_Complex_Float ("C:", C);
Show_Complex_Float ("D:", D);
28
29
       Show_Complex_Float ("X:", X);
30
   end Show_Use;
31
```

Code block metadata

Build output

This example doesn't compile because we have direct visibility to both Complex_Float_Types and Complex_Long_Float_Types packages, and both of them declare the Complex type. In this case, the type declaration becomes hidden, as the compiler cannot decide which declaration of Complex it should take.

13.7.3 Circumventing naming conflicts

As we know, a simple fix for this compilation error is to add the package prefix in the variable declaration:

Listing 74: show_use.adb

```
with Ada.Numerics; use Ada.Numerics;
1
2
   with Ada.Numerics.Generic_Complex_Types;
3
4
   with Show_Any_Complex;
5
6
   procedure Show Use is
7
      package Complex Float Types is new
8
         Ada.Numerics.Generic Complex Types
9
           (Real => Float);
10
      use Complex Float Types;
11
12
      package Complex Long Float Types is new
13
         Ada.Numerics.Generic_Complex_Types
14
           (Real => Long_Float);
15
      use Complex_Long_Float_Types;
16
17
```

```
procedure Show_Complex_Float is new
18
          Show_Any_Complex (Complex_Float_Types);
19
20
       C, D, X : Complex_Float_Types.Complex;
21
                   ^ SOLVED: package is now specified.
22
       - -
   begin
23
       C := Compose_From_Polar (3.0, Pi / 2.0);
24
       D := Compose_From_Polar (5.0, Pi / 2.0);
25
       X := C + D;
26
27
       Show_Complex_Float ("C:", C);
Show_Complex_Float ("D:", D);
28
29
       Show_Complex_Float ("X:", X);
30
   end Show_Use;
31
```

Code block metadata

Runtime output

C: (-1.31134E-07, 3.00000E+00) D: (-2.18557E-07, 5.00000E+00) X: (-3.49691E-07, 8.00000E+00)

Another possibility is to write a use clause in the form **use** all **type**:

Listing 75: show_use.adb

```
with Ada.Numerics; use Ada.Numerics;
1
2
   with Ada.Numerics.Generic_Complex_Types;
3
4
   with Show Any Complex;
5
6
   procedure Show Use is
7
      package Complex Float Types is new
8
         Ada.Numerics.Generic Complex Types
9
           (Real => Float);
10
      use all type Complex_Float_Types.Complex;
11
12
      package Complex_Long_Float_Types is new
13
        Ada.Numerics.Generic_Complex_Types
14
           (Real => Long Float);
15
      use all type Complex_Long_Float_Types.Complex;
16
17
      procedure Show Complex Float is new
18
        Show Any Complex (Complex Float Types);
19
20
      C, D, X : Complex_Float_Types.Complex;
21
   begin
22
      C := Compose_From_Polar (3.0, Pi / 2.0);
23
      D := Compose_From_Polar (5.0, Pi / 2.0);
24
      X := C + D;
25
26
      Show_Complex_Float ("C:", C);
27
      Show_Complex_Float ("D:", D);
28
      Show_Complex_Float ("X:", X);
29
   end Show Use;
30
```

Code block metadata

Runtime output

C: (-1.31134E-07, 3.00000E+00) D: (-2.18557E-07, 5.00000E+00) X: (-3.49691E-07, 8.00000E+00)

For the sake of completeness, let's declare and use variables of both Complex types:

Listing 76: show use.adb

```
with Ada.Numerics; use Ada.Numerics;
1
2
   with Ada.Numerics.Generic_Complex_Types;
3
4
   with Show_Any_Complex;
5
6
   procedure Show Use is
7
      package Complex Float Types is new
8
        Ada.Numerics.Generic_Complex_Types
9
           (Real => Float);
10
      use all type Complex_Float_Types.Complex;
11
12
      package Complex_Long_Float_Types is new
13
         Ada.Numerics.Generic_Complex_Types
14
           (Real => Long_Float);
15
      use all type Complex_Long_Float_Types.Complex;
16
17
      procedure Show Complex Float is new
18
        Show_Any_Complex (Complex_Float_Types);
19
20
      procedure Show Complex Long Float is new
21
         Show_Any_Complex (Complex_Long_Float_Types);
22
23
      C, D, X : Complex_Float_Types.Complex;
24
      E, F, Y : Complex_Long_Float_Types.Complex;
25
26
   begin
      C := Compose_From_Polar (3.0, Pi / 2.0);
27
      D := Compose_From_Polar (5.0, Pi / 2.0);
28
      X := C + D;
29
30
      Show_Complex_Float ("C:", C);
31
      Show_Complex_Float ("D:", D);
32
      Show_Complex_Float ("X:", X);
33
34
      E := Compose_From_Polar (3.0, Pi / 2.0);
35
      F := Compose_From_Polar (5.0, Pi / 2.0);
36
      Y := E + F;
37
38
      Show_Complex_Long_Float ("E:", E);
39
      Show_Complex_Long_Float ("F:", F);
40
      Show_Complex_Long_Float ("Y:", Y);
41
   end Show_Use;
42
```

Code block metadata

Project: Courses.Advanced_Ada.Modular_Prog.Packages.Use_Clause_Naming_Conflicts. □Use_Type_Clause_Complex_Types

MD5: 48f31250116f107d3143703debb3107d

Runtime output

C:	(-1.31134E-07, 3.00000E	+00)
D:	(-2.18557E-07, 5.00000E	+00)
Χ:	(-3.49691E-07, 8.00000E	+00)
E:	(1.83697019872103E-16,	3.00000000000000E+00)
F:	(3.06161699786838E-16,	5.00000000000000E+00)
Υ:	(4.89858719658941E-16,	8.00000000000000E+00)

As expected, the code compiles correctly.

CHAPTER FOURTEEN

SUBPROGRAMS AND MODULARITY

14.1 Private subprograms

We've seen *previously* (page 552) that we can declare private packages. Because packages and subprograms can both be library units, we can declare private subprograms as well. We do this by using the **private** keyword. For example:

Listing 1: test.ads

```
private procedure Test;
1
```

Listing 2: test.adb

procedure Test is 1 begin 2

3

null: end Test;

Code block metadata

```
Project: Courses.Advanced Ada.Modular Prog.Subprograms Modularity.Private
→Subprograms.Private Test Procedure
MD5: 2ea1770a5fd5dee40f015b9d33d2f309
```

Such a subprogram as the one above isn't really useful. For example, we cannot write a with clause that refers to the Test procedure, as it's not visible anywhere:

Listing 3: show test.adb

```
with Test:
1
2
   procedure Show Test is
3
   begin
4
      Test;
5
  end Show Test;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Subprograms_Modularity.Private_
 →Subprograms.Private Test Procedure
MD5: 0702378a034f65a69a4c5b5258f7b32e
```

Build output

```
show_test.adb:1:06: error: current unit must also be private descendant of

→"Standard"
gprbuild: *** compilation phase failed
```

As expected, since Test is private, we get a compilation error because this procedure cannot be referenced in the Show_Test procedure.

1 In the Ada Reference Manual

- 10.1.1 Compilation Units Library Units²⁵⁴
- 10.1.2 Context Clauses With Clauses²⁵⁵

14.1.1 Private subprograms of a package

A more useful example is to declare private subprograms of a package. For example:

```
Listing 4: data_processing.ads
```

```
package Data_Processing is
1
2
      type Data is private;
3
4
      procedure Process (D : in out Data);
5
6
   private
7
8
        type Data is record
9
          F : Float;
10
        end record;
11
12
   end Data_Processing;
13
```

Listing 5: data_processing.adb

```
with Data_Processing.Calculate;
1
2
   package body Data_Processing is
3
4
      procedure Process (D : in out Data) is
5
      begin
6
         Calculate (D);
7
      end Process;
8
9
  end Data_Processing;
10
```

Listing 6: data_processing-calculate.ads

```
private
procedure Data_Processing.Calculate
(D : in out Data);
```

Listing 7: data_processing-calculate.adb

```
1 procedure Data_Processing.Calculate
2 (D : in out Data)
3 is
4 begin
5 -- Dummy implementation...
6 D.F := 0.0;
7 end Data_Processing.Calculate;
```

²⁵⁴ http://www.ada-auth.org/standards/22rm/html/RM-10-1-1.html
 ²⁵⁵ http://www.ada-auth.org/standards/22rm/html/RM-10-1-2.html

Listing 8: test_data_processing.adb

```
with Data_Processing; use Data_Processing;
procedure Test_Data_Processing is
D : Data;
begin
Process (D);
end Test_Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Subprograms_Modularity.Private_

→Subprograms.Private_Package_Procedure

MD5: 0f6af1b02f37e011abac5b2a6dfc482d
```

In this example, we declare Calculate as a private procedure of the Data_Processing package. Therefore, it's visible in that package (but not in the Test_Data_Processing procedure). Also, in the Calculate procedure, we're able to initialize the private component F of the D object because the child subprogram has access to the private part of its parent package.

14.1.2 Private subprograms and private packages

We can also use private subprograms to test private packages. As we know, in most cases, we cannot access private packages in external clients — such as external subprograms. However, by declaring a subprogram private, we're allowed to access private packages. This can be very useful to create applications that we can use to test private packages. (Note that these applications must be library-level parameterless subprograms, because only those can be main programs.)

Let's see an example:

Listing 9: private_data_processing.ads

```
private package Private_Data_Processing is
1
2
       type Data is private;
3
4
      procedure Process (D : in out Data);
5
6
   private
7
8
        type Data is record
9
           F : Float;
10
        end record;
11
12
   end Private Data Processing;
13
```

Listing 10: private_data_processing.adb

```
1 package body Private_Data_Processing is
2
3 procedure Process (D : in out Data) is
4 begin
5 D.F := 0.0;
6 end Process;
7
8 end Private_Data_Processing;
```

Listing 11: test_private_data_processing.ads

```
private procedure Test_Private_Data_Processing;
```

Listing 12: test_private_data_processing.adb

```
1 with Private_Data_Processing;
2 use Private_Data_Processing;
3 
4 procedure Test_Private_Data_Processing is
5 D : Data;
6 begin
7 Process (D);
8 end Test_Private_Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Subprograms_Modularity.Private_

→Subprograms.Private_Subprogram_Private_Package

MD5: 3527e54f99eb2cb52317c987b499caaf
```

In this code example, we have the private **Private**_Data_Processing package. In order to test it, we implement the private procedure Test_Private_Data_Processing. The fact that this procedure is private allows us to use the **Private**_Data_Processing package as if it was a non-private package. We then use the private Test_Private_Data_Processing procedure as our main application, so we can run it to test application the private package.

Child subprograms of private packages

We could also implement the Test subprogram that we use to test a private package P as a child subprogram of that package. In other words, we could write a procedure P.Test and use it as our main application. The advantage here is that this allows us to access the private part of the parent package P in the test procedure.

Let's rewrite the Test_Private_Data_Processing procedure from the previous example as the child procedure **Private**_Data_Processing.Test:

Listing 13: private_data_processing.ads

```
private package Private_Data_Processing is
1
2
       type Data is private;
3
4
      procedure Process (D : in out Data);
5
6
   private
7
8
        type Data is record
9
          F : Float;
10
        end record;
11
12
   end Private_Data_Processing;
13
```

Listing 14: private_data_processing.adb

```
package body Private_Data_Processing is
procedure Process (D : in out Data) is
begin
null:
```

```
6 end Process;
```

8 end Private_Data_Processing;

Listing 15: private_data_processing-test.ads

```
procedure Private_Data_Processing.Test;
```

Listing 16: private_data_processing-test.adb

```
procedure Private_Data_Processing.Test is
D : Data := (F => 0.0);
begin
Process (D);
end Private_Data_Processing.Test;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Modular_Prog.Subprograms_Modularity.Private_

Subprograms.Private_Package_Child_Subprogram

MD5: 0726f5890a5b3847244d1ae08989e158
```

In this code example, we now implement the Test procedure as a child of the **Pri-vate**_Data_Processing package. In this procedure, we're able to initialize the private component F of the D object. As we know, this initialization of a private component wouldn't be possible if Test wasn't a child procedure. (For instance, writing such an initialization in the Test_Private_Data_Processing procedure from the previous code example would trigger a compilation error.)

Part IV

Resource Management

CHAPTER FIFTEEN

ACCESS TYPES

We discussed access types back in the Introduction to Ada course²⁵⁶. In this chapter, we discuss further details about access types and techniques when using them. Before we dig into details, however, we're going to make sure we understand the terminology.

15.1 Access types: Terminology

In this section, we discuss some of the terminology associated with access types. Usually, the terms used in Ada when discussing references and dynamic memory allocation are different than the ones you might encounter in other languages, so it's necessary you understand what each term means.

15.1.1 Access type, designated subtype and profile

The first term we encounter is (obviously) *access type*, which is a type that provides us access to an object or a subprogram. We declare access types by using the **access** keyword:

Listing 1: show_access_type_declaration.ads

```
package Show_Access_Type_Declaration is
1
2
з
          Declaring access types:
      - -
4
5
6
      -- Access-to-object type
7
      type Integer_Access is access Integer;
8
9
      -- Access-to-subprogram type
10
      type Init_Integer_Access is access
11
        function return Integer;
12
13
  end Show Access Type Declaration;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Terminology.Access_

→Type_Declaration

MD5: 64e4e0847a73a9ed23e29e09798934de
```

Here, we're declaring two access types: the access-to-object type Integer_Access and the access-to-subprogram type Init_Integer_Access. (We discuss access-to-subprogram types *later on* (page 677)).

In the declaration of an access type, we always specify — after the **access** keyword — the kind of thing we want to designate. In the case of an access-to-object type declaration,

²⁵⁶ https://learn.adacore.com/courses/intro-to-ada/chapters/access_types.html#intro-ada-access-types-overview

we declare a subtype we want to access, which is known as the *designated subtype* of an access type. In the case of an access-to-subprogram type declaration, the subprogram prototype is known as the *designated profile*.

In our previous code example, **Integer** is the designated subtype of the Integer_Access type, and **function** return Integer is the designated profile of the Init_Integer_Access type.

```
Important
```

In contrast to other programming languages, an access type is not a pointer, and it doesn't just indicate an address in memory. We discuss more about *addresses* (page 706) later on.

15.1.2 Access object and designated object

We use an access-to-object type by first declaring a variable (or constant) of an access type and then allocating an object. (This is actually just one way of using access types; we discuss other methods later in this chapter.) The actual variable or constant of an access type is called *access object*, while the object we allocate (via **new**) is the *designated object*.

For example:

Listing 2: s	show_si	mple_a	allocation.adb
--------------	---------	--------	----------------

```
procedure Show_Simple_Allocation is
1
2
      -- Access-to-object type
3
      type Integer_Access is access Integer;
4
5
      -- Access object
6
      I1 : Integer_Access;
7
8
   begin
9
      I1 := new Integer;
10
             ^^^^ allocating an object,
      - -
11
                         which becomes the designated
      - -
12
                         object for I1
       - -
13
14
  end Show_Simple_Allocation;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Terminology.Simple_

⇔Allocation

MD5: 32ca8cf523e19b25dabb55da6df1f18d
```

In this example, I1 is an access object and the object allocated via **new Integer** is its designated object.

15.1.3 Access value and designated value

An access object and a designated (allocated) object, both store values. The value of an access object is the *access value* and the value of a designated object is the *designated value*. For example:

```
Listing 3: show_values.adb
```

```
procedure Show Values is
1
2
       -- Access-to-object type
3
      type Integer Access is access Integer;
4
5
      I1, I2, I3 : Integer Access;
6
7
   begin
8
      I1 := new Integer;
9
      I3 := new Integer;
10
11
       -- Copying the access value of I1 to I2
12
      I2 := I1:
13
14
       -- Copying the designated value of I1
15
      I3.all := I1.all;
16
17
   end Show Values;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Terminology.Values
MD5: a152ee813b8ed9fad985cf4e2c25d847

In this example, the assignment I2 := I1 copies the access value of I1 to I2. The assignment I3.all := I1.all copies I1's designated value to I3's designated object. (As we already know, .all is used to dereference an access object. We discuss this topic again *later in this chapter* (page 623).)

1 In the Ada Reference Manual

3.10 Access Types²⁵⁷

15.2 Access types: Allocation

Ada makes the distinction between pool-specific and general access types, as we'll discuss in this section. Before doing so, however, let's talk about memory allocation.

In general terms, memory can be allocated dynamically on the heap or statically on the stack. (Strictly speaking, both are dynamic allocations, in that they occur at run-time with amounts not previously specified.) For example:

```
Listing 4: show_simple_allocation.adb
```

```
procedure Show Simple Allocation is
1
2
      -- Declaring access type:
3
      type Integer_Access is access Integer;
4
5
      -- Declaring access object:
6
      A1 : Integer Access;
7
8
   begin
9
          Allocating an Integer object on the heap
10
      - -
```

(continues on next page)

²⁵⁷ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

```
A1 := new Integer;
11
12
       declare
13
         -- Allocating an Integer object on the
14
          -- stack
15
          I : Integer;
16
       begin
17
          null;
18
       end;
19
20
   end Show Simple Allocation;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Types_

⇔Allocation.Simple_Allocation

MD5: 4144feb99e6e0b1a0749fce0b20370a1
```

When we allocate an object on the heap via **new**, the allocation happens in a memory pool that is associated with the access type. In our code example, there's a memory pool associated with the Integer_Access type, and each **new Integer** allocates a new integer object in that pool. Therefore, access types of this kind are called pool-specific access types. (We discuss *more about these types* (page 597) later.)

It is also possible to access objects that were allocated on the stack. To do that, however, we cannot use pool-specific access types because — as the name suggests — they're only allowed to access objects that were allocated in the specific pool associated with the type. Instead, we have to use general access types in this case:

Listing 5: show_general_access_type.adb

```
procedure Show_General_Access_Type is
1
2
      -- Declaring general access type:
3
      type Integer_Access is access all Integer;
4
5
      -- Declaring access object:
6
      A1 : Integer Access;
7
8
      -- Allocating an Integer object on the
9
      -- stack:
10
      I : aliased Integer;
11
12
   beain
13
      -- Getting access to an Integer object that
14
          was allocated on the stack
15
      A1 := I'Access;
16
17
   end Show_General_Access_Type;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Types_

Allocation.General_Access_Types

MD5: f166291ad1975396131775d0aff6ad9d
```

In this example, we declare the general access type Integer_Access and the access object A1. To initialize A1, we write I 'Access to get access to an integer object I that was allocated on the stack. (For the moment, don't worry much about these details: we'll talk about general access types again when we introduce the topic of *aliased objects* (page 636) later on.)

1 For further reading...

Note that it is possible to use general access types to allocate objects on the heap:

Listing 6: show simple allocation.adb

```
procedure Show_Simple_Allocation is
```

```
-- Declaring general access type:
type Integer Access is access all Integer;
```

```
-- Declaring access object:
A1 : Integer_Access;
```

begin

```
JIN
--
--
Allocating an Integer object on the heap
-- and initializing an access object of
-- the general access type Integer_Access.
Al := new Integer;
```

1 2

3

4 5

6

7 8

9 10

11

12

13

end Show_Simple_Allocation;

Code block metadata

Here, we're using a general access type Integer_Access, but allocating an integer object on the heap.

1 Important

In many code examples, we have used the **Integer** type as the designated subtype of the access types — by writing **access Integer**. Although we have used this specific scalar type, we aren't really limited to those types. In fact, we can use *any type* as the designated subtype, including user-defined types, composite types, task types and protected types.

In the Ada Reference Manual

• 3.10 Access Types²⁵⁸

15.2.1 Pool-specific access types

We've already discussed many aspects about pool-specific access types. In this section, we recapitulate some of those aspects, and discuss some new details that haven't seen yet.

As we know, we cannot directly assign an object Distance_Miles of type Miles to an object Distance_Meters of type Meters, even if both share a common **Float** type ancestor. The assignment is only possible if we perform a type conversion from Miles to Meters, or vice-versa — e.g.: Distance_Meters := Meters (Distance_Miles) * Miles_To_Meters_Factor.

²⁵⁸ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

Similarly, in the case of pool-specific access types, a direct assignment between objects of different access types isn't possible. However, even if both access types have the same designated subtype (let's say, they are both declared using **is access Integer**), it's still not possible to perform a type conversion between those access types. The only situation when an access type conversion is allowed is when both types have a common ancestor.

Let's see an example:

Listing 7: show_simple_allocation.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Simple Allocation is
3
4
      -- Declaring pool-specific access type:
5
      type Integer_Access_1 is access Integer;
6
      type Integer_Access_2 is access Integer;
7
      type Integer_Access_2B is new Integer_Access_2;
8
9
      -- Declaring access object:
10
      A1 : Integer_Access_1;
11
      A2 : Integer_Access_2;
12
      A2B : Integer_Access_2B;
13
14
   begin
15
      A1 := new Integer;
16
      Put_Line ("A1 : " & A1'Image);
17
      Put Line ("Pool: " & A1'Storage Pool'Image);
18
19
      A2 := new Integer;
20
      Put Line ("A2: " & A2'Image);
21
      Put_Line ("Pool: " & A2'Storage_Pool'Image);
22
23
          ERROR: Cannot directly assign access values
      - -
24
      - -
                  for objects of unrelated access
25
      - -
                  types; also, cannot convert between
26
       - -
                  these types.
27
       - -
28
      -- A1 := A2;
29
      -- A1 := Integer_Access_1 (A2);
30
31
      A2B := Integer_Access_2B (A2);
32
      Put_Line ("A2B: " & A2B'Image);
33
      Put Line ("Pool: " & A2B'Storage Pool'Image);
34
35
   end Show Simple Allocation;
36
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Types_

⇔Allocation.Pool_Specific_Access_Types

MD5: 80d0e9764917fa8352b6616e3a8425de
```

Runtime output

```
A1 : (access 5895afa1a2a0)

Pool: SYSTEM.POOL_GLOBAL.UNBOUNDED_NO_RECLAIM_POOL'{SYSTEM.STORAGE_POOLS.TROOT_

STORAGE_POOLC object}

A2: (access 5895afa1a360)

Pool: SYSTEM.POOL_GLOBAL.UNBOUNDED_NO_RECLAIM_POOL'{SYSTEM.STORAGE_POOLS.TROOT_

STORAGE_POOLC object}

A2B: (access 5895afa1a360)
```

Pool: SYSTEM.POOL_GLOBAL.UNBOUNDED_NO_RECLAIM_POOL'{SYSTEM.STORAGE_POOLS.TROOT_ STORAGE_POOLC object}

In this example, we declare three access types: Integer Access 1, Integer Access 2 and Integer_Access_2B. Also, the Integer_Access_2B type is derived from the Inte-ger_Access_2 type. Therefore, we can convert an object of Integer_Access_2 type to the Integer_Access_2B type — we do this in the A2B := Integer_Access_2B (A2) assignment. However, we cannot directly assign to or convert between unrelated types such as Integer Access 1 and Integer Access 2. (We would get a compilation error if we included the A1 := A2 or the A1 := Integer Access 1 (A2) assignment.)

```
Important
```

Remember that:

- As mentioned in the Introduction to Ada course²⁵⁹:
 - an access type can be unconstrained, but the actual object allocation must be constrained;
 - we can use a *qualified expression* (page 64) to allocate an object.
- We can use the Storage Size attribute to limit the size of the memory pool associated with an access type, as discussed previously in the section about storage size (page 85).
- When running out of memory while allocating via new, we get a Storage Error exception because of the *storage check* (page 528).

For example:

1

```
Listing 8: show array allocation.adb
```

```
with Ada.Text IO; use Ada.Text IO;
2
   procedure Show_Array_Allocation is
3
4
       -- Unconstrained array type:
5
       type Integer Array is
6
        array (Positive range <>) of Integer;
7
8
       -- Access type with unconstrained
9
       -- designated subtype and limited storage
10
       -- size.
11
       type Integer Array Access is
12
        access Integer Array
13
           with Storage Size => 128;
14
15
       -- An access object:
16
      A1 : Integer Array Access;
17
18
       procedure Show Info
19
        (IAA : Integer_Array_Access) is
20
       begin
21
          Put Line ("Allocated: " & IAA'Image);
22
          Put Line ("Length:
23
                     & IAA.all'Length'Image);
24
          Put_Line ("Values:
25
                     & IAA.all'Image);
26
       end Show Info;
27
28
   begin
29
```

```
-- Allocating an integer array with
30
       -- constrained range on the heap:
31
      A1 := new Integer_Array (1 .. 3);
32
       A1.all := [others => 42];
33
       Show_Info (A1);
34
35
       -- Allocating an integer array on the
36
       -- heap using a qualified expression:
37
      A1 := new Integer_Array'(5, 10);
38
       Show_Info (A1);
39
40
       -- A third allocation fails at run time
41
       -- because of the constrained storage
42
          size:
       - -
43
      A1 := new Integer_Array (1 .. 100);
44
      Show_Info (A1);
45
46
   exception
47
       when Storage Error =>
48
           Put_Line ("Out of memory!");
49
50
   end Show Array Allocation;
51
```

15.2.2 Multiple allocation

Up to now, we have seen examples of allocating a single object on the heap. It's possible to allocate multiple objects *at once* as well — i.e. syntactic sugar is available to simplify the code that performs this allocation. For example:

Listing 9: show_access_array_allocation.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show_Access_Array_Allocation is
3
4
      type Integer_Access is access Integer;
5
6
      type Integer Access Array is
7
        array (Positive range <>) of Integer_Access;
8
9
      -- An array of access objects:
10
      Arr : Integer_Access_Array (1 .. 10);
11
12
   begin
13
14
       -- Allocating 10 access objects and
15
       -- initializing the corresponding designated
16
      -- object with zero:
17
18
      Arr := (others => new Integer'(0));
19
20
       -- Same as:
21
      for I in Arr'Range loop
22
         Arr (I) := new Integer'(0);
23
      end loop;
24
25
      Put_Line ("Arr: " & Arr'Image);
26
27
```

(continues on next page)

²⁵⁹ https://learn.adacore.com/courses/intro-to-ada/chapters/access_types.html#intro-ada-access-type-allocation-constraints

```
28 Put_Line ("Arr (designated values): ");
29 for E of Arr loop
30 Put (E.all'Image);
31 end loop;
32
33 end Show_Access_Array_Allocation;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Types_

⇔Allocation.Integer_Access_Array

MD5: 4afc9358c8aa9426a97ca8932c75d932
```

Runtime output

```
Arr:
[(access 616dcbaaa3e0), (access 616dcbaaa400), (access 616dcbaaa420),
(access 616dcbaaa440), (access 616dcbaaa460), (access 616dcbaaa480),
(access 616dcbaaa4a0), (access 616dcbaaa4c0), (access 616dcbaaa4e0),
(access 616dcbaaa500)]
Arr (designated values):
0 0 0 0 0 0 0 0 0 0 0
```

In this example, we have the access type Integer_Access and an array type of this access type (Integer_Access_Array). We also declare an array Arr of Integer_Access_Array type. This means that each component of Arr is an access object. We allocate all ten components of the Arr array by simply writing Arr := (others => new Integer). This array aggregate (page 262) is syntactic sugar for a loop over Arr that allocates each component. (Note that, by writing Arr := (others => new Integer'(θ)), we're also initializing the designated objects with zero.)

Let's see another code example, this time with task types:

Listing 10: workers.ads

```
package Workers is
1
2
      task type Worker is
3
         entry Start (Id : Positive);
4
          entry Stop;
5
      end Worker;
6
7
      type Worker_Access is access Worker;
8
9
      type Worker Array is
10
         array (Positive range <>) of Worker Access;
11
12
   end Workers;
13
```

Listing 11: workers.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
3
  package body Workers is
4
5
      task body Worker is
         Id : Positive;
6
      begin
7
         accept Start (Id : Positive) do
8
            Worker.Id := Id;
9
```

```
end Start;
10
          Put_Line ("Started Worker #"
11
                      & Id'Image);
12
13
          accept Stop;
14
15
          Put_Line ("Stopped Worker #"
16
                      & Id'Image);
17
       end Worker;
18
19
   end Workers;
20
```

Listing 12: show_workers.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Workers; use Workers;
3
4
   procedure Show_Workers is
5
      Worker_Arr : Worker_Array (1 .. 20);
6
   begin
7
8
       - -
       -- Allocating 20 workers at once:
9
10
      Worker_Arr := (others => new Worker);
11
12
       for I in Worker_Arr'Range loop
13
          Worker_Arr (I).Start (I);
14
15
      end loop;
16
      Put_Line ("Some processing...");
17
      delay 1.0;
18
19
      for W of Worker_Arr loop
20
          W.Stop;
21
      end loop;
22
23
   end Show Workers;
24
```

Code block metadata

Runtime output

Started Worker # 3 Started Worker # 4 Started Worker # 2 Started Worker # 1 Started Worker # 5 Started Worker # 6 Started Worker # 7 Started Worker # 8 Started Worker # 10 Started Worker # 11 Started Worker # 11 Started Worker # 13 Started Worker # 14

Started	Worker # 15
Started	Worker # 16
Started	Worker # 17
Started	Worker # 18
Started	Worker # 19
Started	Worker # 20
Some pro	cessing
Stopped	Worker # 5
Stopped	Worker # 1
Stopped	Worker # 6
Stopped	Worker # 7
Stopped	Worker # 8
Stopped	Worker # 9
Stopped	Worker # 10
Stopped	Worker # 11
Stopped	Worker # 12
Stopped	Worker # 13
Stopped	Worker # 14
Stopped	Worker # 15
Stopped	Worker # 16
Stopped	Worker # 17
Stopped	Worker # 18
Stopped	Worker # 19
Stopped	Worker # 20
Stopped	Worker # 3
Stopped	Worker # 2
Stopped	Worker # 4
-	

In this example, we declare the task type Worker, the access type Worker_Access and an array of access to tasks Worker_Array. Using this approach, a task is only created when we allocate an individual component of an array of Worker_Array type. Thus, when we declare the Worker_Arr array in this example, we're only preparing a *container* of 20 workers, but we don't have any actual tasks yet. We bring the 20 tasks into existence by writing Worker Arr := (**others** => **new** Worker).

15.3 Discriminants as Access Values

We can use access types when declaring discriminants. Let's see an example:

```
Listing 13: custom_recs.ads
```

```
2
      -- Declaring an access type:
3
      type Integer_Access is access Integer;
4
5
      -- Declaring a discriminant with this
6
      -- access type:
7
      type Rec (IA : Integer_Access) is record
8
9
         I : Integer := IA.all;
10
                      ~~~~~~
11
          -- Setting I's default to use the
12
         -- designated value of IA:
13
      end record;
14
15
      procedure Show (R : Rec);
16
17
  end Custom_Recs;
18
```

package Custom Recs is

1

Listing 14: custom_recs.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Custom Recs is
3
4
      procedure Show (R : Rec) is
5
      begin
6
          Put_Line ("R.IA = "
7
                    & Integer'Image (R.IA.all));
8
          Put_Line ("R.I = "
9
                    & Integer'Image (R.I));
10
      end Show;
11
12
   end Custom Recs;
13
```

Listing 15: show_discriminants_as_access_values.adb

```
with Custom_Recs; use Custom_Recs;
1
2
   procedure Show_Discriminants_As_Access_Values is
3
4
      IA : constant Integer_Access :=
5
              new Integer'(10);
6
      R : Rec (IA);
7
8
9
   begin
      Show (R);
10
11
      IA.all := 20;
12
      R.I := 30;
13
      Show (R);
14
15
       -- As expected, we cannot change the
16
          discriminant. The following line is
      - -
17
          triggers a compilation error:
      - -
18
19
       -- R.IA := new Integer;
20
21
   end Show_Discriminants_As_Access_Values;
22
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Discriminants_As_ ⇔Access_Values.Discriminant_Access_Values MD5: c7850acefd8e5227f4be654faed13055

Runtime output

R.IA = 10 R.I = 10 R.IA = 20 R.I = 30

In the Custom_Recs package from this example, we declare the access type Integer_Access. We then use this type to declare the discriminant (IA) of the Rec type. In the Show_Discriminants_As_Access_Values procedure, we see that (as expected) we cannot change the discriminant of an object of Rec type: an assignment such as R.IA := **new Integer** would trigger a compilation error.

Note that we can use a default for the discriminant:

Listing 16: custom_recs.ads

```
package Custom_Recs is
1
2
      type Integer_Access is access Integer;
3
4
      type Rec (IA : Integer Access
5
                        := new Integer'(0)) is
6
                             ~~~~~~
          - -
7
         - -
                             default value
8
      record
9
       I : Integer := IA.all;
10
      end record;
11
12
      procedure Show (R : Rec);
13
14
   end Custom Recs;
15
```

Listing 17: show_discriminants_as_access_values.adb

```
with Custom_Recs; use Custom_Recs;
1
2
   procedure Show_Discriminants_As_Access_Values is
3
4
5
      R1 : Rec;
           ~~~
      - -
6
          no discriminant: use default
7
      - -
8
      R2 : Rec (new Integer'(20));
9
                 ~~~~~~
10
      - -
                 allocating an unnamed integer object
11
12
   begin
13
      Show (R1);
14
      Show (R2);
15
   end Show_Discriminants_As_Access_Values;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Discriminants_As_
⇔Access_Values.Discriminant_Access_Values
MD5: 968cb88ed7e9e6958ab66fb6f5a7ce2d
```

Runtime output

R.IA = 0 R.I = 0 R.IA = 20 R.I = 20

Here, we've changed the declaration of the Rec type to allocate an integer object if the type's discriminant isn't provided — we can see this in the declaration of the R1 object in the Show_Discriminants_As_Access_Values procedure. Also, in this procedure, we're allocating an unnamed integer object in the declaration of R2.

() In the Ada Reference Manual

- 3.10 Access Types²⁶⁰
- 3.7.1 Discriminant Constraints²⁶¹

15.3.1 Unconstrained type as designated subtype

Notice that we were using a scalar type as the designated subtype of the Integer_Access type. We could have used an unconstrained type as well. In fact, this is often used for the sake of having the effect of an unconstrained discriminant type.

Let's see an example:

Listing 18: persons.ads

```
package Persons is
1
2
      -- Declaring an access type whose
3
      -- designated subtype is unconstrained:
4
      type String Access is access String;
5
6
      -- Declaring a discriminant with this
7
      -- access type:
8
      type Person (Name : String_Access) is record
9
         Age : Integer;
10
      end record;
11
12
      procedure Show (P : Person);
13
14
   end Persons;
15
```

Listing 19: persons.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Persons is
3
4
      procedure Show (P : Person) is
5
      begin
6
          Put Line ("Name = "
7
                    & P.Name.all);
8
          Put_Line ("Age = "
9
                    & Integer'Image (P.Age));
10
      end Show;
11
12
   end Persons;
13
```

Listing 20: show_person.adb

```
with Persons; use Persons;
procedure Show_Person is
P : Person (new String'("John"));
begin
P.Age := 30;
Show (P);
end Show_Person;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Discriminants_As_ ⊶Access_Values.Persons MD5: 9b1109d076b6f06632c8685a41616210

Runtime output

²⁶⁰ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html
²⁶¹ http://www.ada-auth.org/standards/22rm/html/RM-3-7-1.html

Name	=	John
Age	=	30

In this example, the discriminant of the Person type has an unconstrained designated type. In the Show_Person procedure, we declare the P object and specify the constraints of the allocated string object — in this case, a four-character string initialized with the name "John".

```
• For further reading...
```

In the previous code example, we used an array — actually, a string — to demonstrate the advantage of using discriminants as access values, for we can use an unconstrained type as the designated subtype. In fact, as we discussed *earlier in another chapter* (page 196), we can only use discrete types (or access types) as discriminants. Therefore, you wouldn't be able to use a string, for example, directly as a discriminant without using access types:

Listing 21: persons.ads

8

۵

```
package Persons is
    -- ERROR: Declaring a discriminant with an
    -- unconstrained type:
    type Person (Name : String) is record
        Age : Integer;
    end record;
```

end Persons;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Discriminants_As_ ⇔Access_Values.Persons_Error MD5: 4144852aaf95da62bc4781b1e8dc2717

Build output

```
persons.ads:5:24: error: discriminants must have a discrete or access type
gprbuild: *** compilation phase failed
```

As expected, compilation fails for this code because the discriminant of the Person type is indefinite.

However, the advantage of discriminants as access values isn't restricted to being able to use unconstrained types such as arrays: we could really use any type as the designated subtype! In fact, we can generalized this to:

```
Listing 22: gen_custom_recs.ads
```

```
generic
1
      type T (<>); -- any type
2
      type T Access is access T;
3
  package Gen_Custom_Recs is
4
      -- Declare a type whose discriminant D can
5
      -- access any type:
6
      type T_Rec (D : T_Access) is null record;
7
  end Gen_Custom_Recs;
8
```

```
Listing 23: custom recs.ads
   with Gen Custom Recs;
1
2
   package Custom_Recs is
3
4
        type Incomp;
5
        -- Incomplete type declaration!
6
7
       type Incomp Access is access Incomp;
8
9
       -- Instantiating package using
10
       -- incomplete type Incomp:
11
      package Inst is new
12
         Gen_Custom_Recs
13
                => Incomp,
           (T
14
            T Access => Incomp_Access);
15
       subtype Rec is Inst.T Rec;
16
17
       -- At this point, Rec (Inst.T Rec) uses
18
          an incomplete type as the designated
       - -
19
          subtype of its discriminant type
       - -
20
21
       procedure Show (R : Rec) is null;
22
23
       -- Now, we complete the Incomp type:
24
      type Incomp (B : Boolean := True) is private;
25
26
   private
27
      -- Finally, we have the full view of the
28
       -- Incomp type:
29
       type Incomp (B : Boolean := True) is
30
        null record;
31
32
   end Custom Recs;
33
                                  Listing 24: show rec.adb
   with Custom Recs; use Custom Recs;
1
2
    procedure Show Rec is
3
      R : Rec (new Incomp);
4
```

5

7

begin Show (R);

end Show_Rec;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Discriminants_As_ ⇔Access_Values.Generic_Access MD5: c65510e8c6a7625cbd08aa9e68f05115

In the Gen_Custom_Recs package, we're using **type T** (<>) — which can be any type — for the designated subtype of the access type T_Access, which is the type of T_Rec's discriminant. In the Custom_Recs package, we use the incomplete type Incomp to instantiate the generic package. Only after the instantiation, we declare the complete type.

Later on, we'll discuss discriminants again when we look into *anonymous access discriminants* (page 725), which provide some advantages in terms of *accessibility rules* (page 645).

15.3.2 Whole object assignments

As expected, we cannot change the discriminant value in whole object assignments. If we do that, the Constraint_Error exception is raised at runtime:

```
Listing 25: show_person.adb
```

```
with Persons; use Persons;
1
2
   procedure Show Person is
3
      S1 : String_Access := new String'("John");
4
      S2 : String Access := new String'("Mark");
5
      P : Person := (Name => S1,
6
                     Age => 30);
7
   begin
8
      P := (Name => S1, Age => 31);
9
                    ^^ OK: we didn't change the
      - -
10
                           discriminant.
11
      Show (P);
12
13
      -- We can just repeat the discriminant:
14
      P := (Name => P.Name, Age => 32);
15
                   ^^^^^ OK: we didn't change the
      - -
16
      - -
                               discriminant.
17
      Show (P);
18
19
      -- Of course, we can change the string itself:
20
      S1.all := "Mark";
21
      Show (P);
22
23
      24
25
                              discriminant!
26
      Show (P);
27
   end Show_Person;
28
```

Code block metadata

Runtime output

Name = John Age = 31 Name = John Age = 32 Name = Mark Age = 32

raised CONSTRAINT_ERROR : show_person.adb:24 discriminant check failed

The first and the second assignments to P are OK because we didn't change the discriminant. However, the last assignment raises the Constraint_Error exception at runtime because we're changing the discriminant.

15.4 Parameters as Access Values

In addition to *using discriminants as access values* (page 603), we can use access types for subprogram formal parameters. For example, the N parameter of the Show procedure below has an access type:

Listing 26: names.ads

```
1 package Names is
2
3 type Name is access String;
4
5 procedure Show (N : Name);
6
7 end Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⇔Access_Values.Names

MD5: 82ce94987dce9026aed54a0deb3cc548
```

This is the complete code example:

Listing 27: names.ads

```
1 package Names is
2
3 type Name is access String;
4
5 procedure Show (N : Name);
6
7 end Names;
```

Listing 28: names.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Names is
3
4
      procedure Show (N : Name) is
5
      begin
6
         Put_Line ("Name: " & N.all);
7
      end Show;
8
9
   end Names;
10
```

Listing 29: show_names.adb

```
with Names; use Names;
procedure Show_Names is
N : Name := new String'("John");
begin
Show (N);
r end Show_Names;
```

Code block metadata

MD5: 526baf1996b4a2970c3fa2e3485dcbad

Runtime output

Name: John

Note that in this example, the Show procedure is basically just displaying the string. Since the procedure isn't doing anything that justifies the need for an access type, we could have implemented it with a *simpler* type:

Listing 30: names.ads

```
1 package Names is
2
3 type Name is access String;
4
5 procedure Show (N : String);
6
7 end Names;
```

Listing 31: names.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Names is
3
4
      procedure Show (N : String) is
5
      begin
6
          Put_Line ("Name: " & N);
7
      end Show;
8
9
   end Names;
10
```

Listing 32: show_names.adb

```
with Names; use Names;
procedure Show_Names is
N : Name := new String'("John");
begin
Show (N.all);
r end Show Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⇔Access_Values.Names_String

MD5: 097ec1ff781fda9deed1de23cae39ae5
```

Runtime output

Name: John

It's important to highlight the difference between passing an access value to a subprogram and passing an object by reference. In both versions of this code example, the compiler will make use of a reference for the actual parameter of the N parameter of the Show procedure. However, the difference between these two cases is that:

• N : Name is a reference to an object (because it's an access value) that is passed by value, and

• N : **String** is an object passed by reference.

15.4.1 Changing the referenced object

Since the Name type gives us access to an object in the Show procedure, we could actually change this object inside the procedure. To illustrate this, let's change the Show procedure to lower each character of the string before displaying it (and rename the procedure to Lower_And_Show):

Listing 33: names.ads

```
1 package Names is
2
3 type Name is access String;
4
5 procedure Lower_And_Show (N : Name);
6
7 end Names;
```

Listing 34: names.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Characters.Handling;
3
   use Ada.Characters.Handling;
4
5
   package body Names is
6
7
      procedure Lower And Show (N : Name) is
8
      begin
9
          for I in N'Range loop
10
             N (I) := To_Lower (N (I));
11
         end loop;
12
         Put_Line ("Name: " & N.all);
13
      end Lower_And_Show;
14
15
   end Names;
16
```

Listing 35: show_changed_names.adb

```
with Names; use Names;
procedure Show_Changed_Names is
N : Name := new String'("John");
begin
Lower_And_Show (N);
end Show_Changed_Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⊶Access_Values.Changed_Names

MD5: 063a507284f5e7ffa669db2c8fdd3d6f
```

Runtime output

Name: john

Notice that, again, we could have implemented the Lower_And_Show procedure without using an access type:

Listing 36: names.ads

```
package Names is
type Name is access String;
procedure Lower_And_Show (N : in out String);
end Names;
```

Listing 37: names.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Characters.Handling;
3
   use Ada.Characters.Handling;
4
5
   package body Names is
6
7
8
      procedure Lower_And_Show (N : in out String) is
9
      begin
          for I in N'Range loop
10
             N (I) := To_Lower (N (I));
11
          end loop;
12
         Put_Line ("Name: " & N);
13
      end Lower_And_Show;
14
15
   end Names;
16
```

Listing 38: show_changed_names.adb

```
with Names; use Names;
procedure Show_Changed_Names is
N : Name := new String'("John");
begin
Lower_And_Show (N.all);
end Show_Changed_Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⇔Access_Values.Changed_Names_String

MD5: 783ea8c45ed8ad3e0007524c11b6bcc4
```

Runtime output

Name: john

15.4.2 Replace the access value

Instead of changing the object in the Lower_And_Show procedure, we could replace the access value by another one — for example, by allocating a new string inside the procedure. In this case, we have to pass the access value by reference using the **in out** parameter mode:

Listing 39: names.ads

1	package	Names	is
2			

```
3 type Name is access String;
4 
5 procedure Lower_And_Show (N : in out Name);
6 
7 end Names;
```

Listing 40: names.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Ada.Characters.Handling;
3
   use Ada.Characters.Handling;
4
5
   package body Names is
6
7
      procedure Lower And Show (N : in out Name) is
8
      begin
9
         N := new String'(To Lower (N.all));
10
         Put Line ("Name: " & N.all);
11
      end Lower And Show;
12
13
   end Names;
14
```

Listing 41: show_changed_names.adb

```
with Names; use Names;
procedure Show_Changed_Names is
N : Name := new String'("John");
begin
Lower_And_Show (N);
r end Show_Changed_Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⇔Access_Values.Replaced_Names

MD5: a4abfe6fdb1e5029e8eea17641cd960b
```

Runtime output

Name: john

Now, instead of changing the object referenced by N, we're actually replacing it with a new object that we allocate inside the Lower_And_Show procedure.

As expected, contrary to the previous examples, we cannot implement this code by relying on parameter modes to replace the object. In fact, we have to use access types for this kind of operations.

Note that this implementation creates a memory leak. In a proper implementation, we should make sure to *deallocate the object* (page 656), as explained later on.

15.4.3 Side-effects on designated objects

In previous code examples from this section, we've seen that passing a parameter by reference using the **in** or **in out** parameter modes is an alternative to using access values as parameters. Let's focus on the subprogram declarations of those code examples and their parameter modes:

Subprogram	Parameter type	Parameter mode	
Show	Name	in	
Show	String	in	
Lower_And_Show	Name	in	
Lower_And_Show	String	in out	

When we analyze the information from this table, we see that in the case of using strings with different parameter modes, we have a clear indication whether the subprogram might change the object or not. For example, we know that a call to Show (N : **String**) won't change the string object that we're passing as the actual parameter.

In the case of passing an access value, we cannot know whether the designated object is going to be altered by a call to the subprogram. In fact, in both Show and Lower_And_Show procedures, the parameter is the same: N : Name — in other words, the parameter mode is **in** in both cases. Here, there's no clear indication about the effects of a subprogram call on the designated object.

The simplest way to ensure that the object isn't changed in the subprogram is by using *access-to-constant types* (page 637), which we discuss later on. In this case, we're basically saying that the object we're accessing in Show is constant, so we cannot possibly change it:

Listing 42: names.ads

```
1 package Names is
2
3 type Name is access String;
4
5 type Constant_Name is access constant String;
6
7 procedure Show (N : Constant_Name);
8
9 end Names;
```

Listing 43: names.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   -- with Ada.Characters.Handling;
3
   -- use Ada.Characters.Handling;
4
5
   package body Names is
6
7
      procedure Show (N : Constant_Name) is
8
      begin
9
             for I in N'Range loop
         - -
10
                N(I) := To\_Lower(N(I));
          - -
11
          -- end loop;
12
         Put_Line ("Name: " & N.all);
13
      end Show;
14
15
```

```
16 end Names;
```

Listing 44: show_names.adb

```
with Names; use Names;
procedure Show_Names is
N : Name := new String'("John");
```

5 begin

- 6 Show (Constant_Name (N));
- 7 end Show_Names;

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⇔Access_Values.Names_Constant

MD5: 77526e0a159bf1bcbef08a21be250f3c
```

Runtime output

Name: John

In this case, the Constant_Name type ensures that the N parameter won't be changed in the Show procedure. Note that we need to convert from Name to Constant_Name to be able to call the Show procedure (in the Show_Names procedure). Although using **in String** is still a simpler solution, this approach works fine.

(Feel free to uncomment the call to To_Lower in the Show procedure and the corresponding with- and use-clauses to see that the compilation fails when trying to change the constant object.)

We could also mitigate the problem by using contracts. For example:

Listing 45: names.ads

```
package Names is
1
2
      type Name is access String;
3
4
      procedure Show (N : Name)
5
       with Post => N.all'Old = N.all;
6
                      ~~~~~
7
             we promise that we won't change
      - -
8
      - -
              the object
9
10
   end Names;
11
```

Listing 46: names.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
       with Ada.Characters.Handling;
3
       use Ada.Characters.Handling;
   - -
4
5
   package body Names is
6
7
      procedure Show (N : Name) is
8
      begin
9
         -- for I in N'Range loop
10
                N (I) := To Lower (N (I));
          - -
11
         -- end loop;
12
         Put_Line ("Name: " & N.all);
13
      end Show;
14
15
   end Names;
16
```

Listing 47: show_names.adb

```
with Names; use Names;
procedure Show_Names is
N : Name := new String'("John");
begin
Show (N);
r end Show_Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_

⇔Access_Values.Names_Postcondition

MD5: 2a70993232baca9d58d36e537a6fd32b
```

Runtime output

Name: John

Although a bit more verbose than a simple **in String**, the information in the specification of Show at least gives us an indication that the object won't be affected by the call to this subprogram. Note that this code actually compiles if we try to modify **N.all** in the Show procedure, but the post-condition fails at runtime when we do that.

(By uncommentating and building the code again, you'll see an exception being raised at runtime when trying to change the object.)

In the postcondition above, we're using '0ld to refer to the original object before the subprogram call. Unfortunately, we cannot use this attribute when dealing with *limited private types* (page 787) — or limited types in general. For example, let's change the declaration of Name and have it as a limited private type instead:

Listing 48: names.ads

```
package Names is
1
2
       type Name is limited private;
3
4
       function Init (S : String) return Name;
5
6
7
       function Equal (N1, N2 : Name)
8
                        return Boolean;
9
       procedure Show (N : Name)
10
         with Post => Equal (N'Old = N);
11
12
   private
13
14
       type Name is access String;
15
16
       function Init (S : String) return Name is
17
         (new String'(S));
18
19
20
       function Equal (N1, N2 : Name)
                        return Boolean is
21
         (N1.all = N2.all);
22
23
   end Names;
24
```

Listing 49: names.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   -- with Ada.Characters.Handling;
3
   -- use Ada.Characters.Handling;
4
5
   package body Names is
6
7
      procedure Show (N : Name) is
8
      begin
9
             for I in N'Range loop
         - -
10
          - -
                N(I) := To\_Lower(N(I));
11
          -- end loop;
12
         Put Line ("Name: " & N.all);
13
      end Show;
14
15
  end Names;
16
```

Listing 50: show_names.adb

```
with Names; use Names;
procedure Show_Names is
N : Name := Init ("John");
begin
Show (N);
r end Show_Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Parameters_As_
⇔Access_Values.Names_Limited_Private
MD5: 39691394d7a934869dc569eb72d1bf3a
```

Build output

```
names.ads:11:26: error: attribute "Old" cannot apply to limited objects
gprbuild: *** compilation phase failed
```

In this case, we have no means to indicate that a call to Show won't change the internal state of the actual parameter.

For further reading...

As an alternative, we could declare a new Constant_Name type that is also limited private. If we use this type in Show procedure, we're at least indicating (in the type name) that the type is supposed to be constant — even though we're not directly providing means to actually ensure that no modifications occur in a call to the procedure. However, the fact that we declare this type as an access-to-constant (in the private part of the specification) makes it clear that a call to Show won't change the designated object.

Let's look at the adapted code:

Listing 51: names.ads

```
package Names is
type Name is limited private;
```

```
type Constant_Name is limited private;
```

5

```
6
       function Init (S : String) return Name;
7
8
       function To_Constant_Name
9
         (N : Name)
10
          return Constant_Name;
11
12
       procedure Show (N : Constant_Name);
13
14
   private
15
16
       type Name is
17
        access String;
18
19
      type Constant_Name is
20
        access constant String;
21
22
       function Init (S : String) return Name is
23
         (new String'(S));
24
25
      function To_Constant_Name
26
         (N : Name)
27
          return Constant_Name is
28
            (Constant_Name (N));
29
30
   end Names;
31
                                   Listing 52: names.adb
   with Ada.Text_I0; use Ada.Text_I0;
1
2
    -- with Ada.Characters.Handling;
3
   -- use Ada.Characters.Handling;
4
5
   package body Names is
6
7
      procedure Show (N : Constant_Name) is
8
      begin
9
         -- for I in N'Range loop
10
          - -
                N (I) := To_Lower (N (I));
11
          -- end loop;
12
          Put Line ("Name: " & N.all);
13
14
      end Show;
15
   end Names;
16
                                 Listing 53: show_names.adb
   with Names; use Names;
1
2
   procedure Show Names is
3
      N : Name := Init ("John");
4
   begin
5
      Show (To Constant Name (N));
6
   end Show Names;
    Code block metadata
    Project: Courses.Advanced Ada.Resource Management.Access Types.Parameters As
     →Access_Values.Names_Constant_Limited_Private
   MD5: 30da588b57e6b4dfbf9934f77d348473
    Runtime output
```

Name: John

In this version of the source code, the Show procedure doesn't have any side-effects, as we cannot modify N inside the procedure.

Having the information about the effects of a subprogram call to an object is very important: we can use this information to set expectations — and avoid unexpected changes to an object. Also, this information can be used to prove that a program works as expected. Therefore, whenever possible, we should avoid access values as parameters. Instead, we can rely on appropriate parameter modes and pass an object by reference.

There are cases, however, where the design of our application doesn't permit replacing the access type with simple parameter modes. Whenever we have an abstract data type encapsulated as a limited private type — such as in the last code example —, we might have no means to avoid access values as parameters. In this case, using the access type is of course justifiable. We'll see such a case in the *next section* (page 620).

15.5 Self-reference

As we've discussed in the section about incomplete types <Adv_Ada_Incomplete_Types>, we can use incomplete types to create a recursive, self-referencing type. Let's revisit a code example from that section:

Listing 54: linked_list_example.ads

```
package Linked List Example is
1
2
      type Integer List;
3
4
      type Next is access Integer List;
5
6
      type Integer_List is record
7
         I : Integer;
8
         N : Next;
9
      end record;
10
11
  end Linked_List_Example;
12
```

Code block metadata

Here, we're using the incomplete type Integer_List in the declaration of the Next type, which we then use in the complete declaration of the Integer_List type.

Self-references are useful, for example, to create unbounded containers — such as the linked lists mentioned in the example above. Let's extend this code example and partially implement a generic package for linked lists:

Listing 55: linked lists.ads

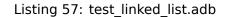
```
1 generic
2 type T is private;
3 package Linked_Lists is
4
5 type List is limited private;
6
7 procedure Append_Front
8 (L : in out List;
```

```
Ε:
                        T);
9
10
       procedure Append_Rear
11
          (L : in out List;
12
           Ε:
                       T);
13
14
       procedure Show (L : List);
15
16
   private
17
18
       -- Incomplete type declaration:
19
       type Component;
20
21
       -- Using incomplete type:
22
       type List is access Component;
23
24
       type Component is record
25
          Value : T;
26
          Next : List;
27
                   ~~~
28
          - -
                Self-reference via access type
29
       end record;
30
31
   end Linked_Lists;
32
```



```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Linked_Lists is
3
4
       procedure Append_Front
5
          (L : in out List;
6
           Ε:
                      T)
7
       is
8
          New First : constant List := new
9
            Component'(Value => E,
10
                        Next => L);
11
       begin
12
          L := New_First;
13
       end Append_Front;
14
15
       procedure Append_Rear
16
          (L : in out List;
17
           Ε:
                       T)
18
       is
19
          New Last : constant List := new
20
            Component'(Value => E,
21
                        Next => null);
22
       begin
23
          if L = null then
24
             L := New_Last;
25
          else
26
             declare
27
                Last : List := L;
28
             begin
29
                while Last.Next /= null loop
30
                    Last := Last.Next;
31
                 end loop;
32
                 Last.Next := New_Last;
33
             end;
34
```

```
end if;
35
       end Append_Rear;
36
37
       procedure Show (L : List) is
38
          Curr : List := L;
39
       begin
40
          if L = null then
41
              Put_Line ("[ ]");
42
          else
43
              Put ("[");
44
              loop
45
                 Put (Curr.Value'Image);
46
                 Put (" ");
47
                 exit when Curr.Next = null;
48
                 Curr := Curr.Next;
49
              end loop;
50
              Put_Line ("]");
51
          end if;
52
       end Show;
53
54
   end Linked_Lists;
55
```



```
with Linked Lists;
1
2
   procedure Test Linked List is
3
        package Integer Lists is new
4
          Linked_Lists (T => Integer);
5
        use Integer_Lists;
6
7
        L : List;
8
   begin
9
        Append_Front (L, 3);
10
        Append_Rear (L, 4);
11
        Append Rear (L, 5);
12
        Append_Front (L, 2);
13
        Append_Front (L, 1);
14
        Append Rear (L, 6);
15
        Append_Rear (L, 7);
16
17
18
        Show (L);
   end Test_Linked_List;
19
```

Code block metadata

Runtime output

[1234567]

In this example, we declare an incomplete type Component in the private part of the generic Linked_Lists package. We use this incomplete type to declare the access type List, which is then used as a self-reference in the Next component of the Component type.

Note that we're using the List type *as a parameter* (page 610) for the Append_Front, Append_Rear and Show procedures.

```
1 In the Ada Reference Manual
```

• 3.10.1 Incomplete Type Declarations²⁶²

15.6 Mutually dependent types using access types

In the section on *mutually dependent types* (page 177), we've seen a code example where each type depends on the other one. We could rewrite that code example using access types:

Listing 58: mutually dependent.ads

```
package Mutually_Dependent is
1
2
       type T2;
3
       type T2_Access is access T2;
4
5
       type T1 is record
6
          B : T2 Access;
7
       end record;
8
9
       type T1 Access is access T1;
10
11
       type T2 is record
12
          A : T1_Access;
13
       end recor\overline{d};
14
15
   end Mutually Dependent;
16
```

Code block metadata

In this example, T1 and T2 are mutually dependent types via the access types T1_Access and T2_Access — we're using those access types in the declaration of the B and A components.

15.7 Dereferencing

In the Introduction to Ada course²⁶³, we discussed the **.all** syntax to dereference access values:

Listing	59:	show	derefer	encing.a	adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Show_Dereferencing is
-- Declaring access type:
type Integer_Access is access Integer;
-- Declaring access object:
```

(continues on next page)

²⁶² http://www.ada-auth.org/standards/22rm/html/RM-3-10-1.html

²⁶³ https://learn.adacore.com/courses/intro-to-ada/chapters/access_types.html#intro-ada-access-dereferencing

```
A1 : Integer_Access;
9
10
   begin
11
      A1 := new Integer;
12
13
       -- Dereferencing access value:
14
      A1.all := 22;
15
16
      Put_Line ("A1: " & Integer'Image (A1.all));
17
   end Show_Dereferencing;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Dereferencing.

→Simple_Dereferencing

MD5: 65655768c17a02991ffeda9a853b6ffb
```

Runtime output

A1: 22

In this example, we declare A1 as an access object, which allows us to access objects of **Integer** type. We dereference A1 by writing A1.all.

Here's another example, this time with an array:

```
Listing 60: show_dereferencing.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Dereferencing is
3
4
      type Integer_Array is
5
        array (Positive range <>) of Integer;
6
7
      type Integer_Array_Access is
8
        access Integer_Array;
9
10
      Arr : constant Integer_Array_Access :=
11
                        new Integer_Array (1 .. 6);
12
   begin
13
      Arr.all := (1, 2, 3, 5, 8, 13);
14
15
      for I in Arr'Range loop
16
         Put_Line ("Arr (: '
17
                    & Integer'Image (I) & "): "
18
                    & Integer'Image (Arr.all (I)));
19
      end loop;
20
   end Show Dereferencing;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Dereferencing.Array_

⊸Dereferencing

MD5: 0e533dfd8ec1a74af17c99633c292e95
```

Runtime output

Arr (: 1): 1 Arr (: 2): 2 Arr (: 3): 3

Arr (: 4): 5 Arr (: 5): 8 Arr (: 6): 13

In this example, we dereference the access value by writing Arr.all. We then assign an array aggregate to it — this becomes Arr.all := (..., ...);. Similarly, in the loop, we write Arr.all (I) to access the I component of the array.

1 In the Ada Reference Manual

• 4.1 Names²⁶⁴

15.7.1 Implicit Dereferencing

Implicit dereferencing allows us to omit the .**all** suffix without getting a compilation error. In this case, the compiler *knows* that the dereferenced object is implied, not the access value.

Ada supports implicit dereferencing in these use cases:

- when accessing components of a record or an array including array slices.
- when accessing subprograms that have at least one parameter (we discuss this topic later in this chapter);
- when accessing some attributes such as some array and task attributes.

Arrays

Let's start by looking into an example of implicit dereferencing of arrays. We can take the previous code example and replace Arr.all (I) by Arr (I):

Listing 61: show_dereferencing.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_Dereferencing is
3
4
      type Integer_Array is
5
        array (Positive range <>) of Integer;
6
7
      type Integer_Array_Access is
8
         access Integer_Array;
9
10
      Arr : constant Integer_Array_Access :=
11
                         new Integer_Array (1 .. 6);
12
   begin
13
      Arr.all := (1, 2, 3, 5, 8, 13);
14
15
      Arr (1 .. 6) := (1, 2, 3, 5, 8, 13);
16
17
      for I in Arr'Range loop
18
          Put_Line
19
            ("Arr (: "
20
             & Integer'Image (I) & "): "
21
             & Integer'Image (Arr (I)));
22
                                   ^ .all is implicit.
23
```

(continues on next page)

²⁶⁴ http://www.ada-auth.org/standards/22rm/html/RM-4-1.html

24 end loop;

25 end Show_Dereferencing;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Dereferencing.Array_ ⊲Implicit_Dereferencing MD5: ade602a9e6976018e0c00f930a2399f1

Runtime output

Arr (: 1): 1 Arr (: 2): 2 Arr (: 3): 3 Arr (: 4): 5 Arr (: 5): 8 Arr (: 6): 13

Both forms — Arr.all (I) and Arr (I) — are equivalent. Note, however, that there's no implicit dereferencing when we want to access the whole array. (Therefore, we cannot write Arr := (1, 2, 3, 5, 8, 13);) However, as slices are implicitly dereferenced, we can write Arr (1 ... 6) := (1, 2, 3, 5, 8, 13); instead of Arr.all (1 ... 6) := (1, 2, 3, 5, 8, 13); Alternatively, we can assign to the array components individually and use implicit dereferencing for each component:

Arr (1) := 1; Arr (2) := 2; Arr (3) := 3; Arr (4) := 5; Arr (5) := 8; Arr (6) := 13;

Implicit dereferencing isn't available for the whole array because we have to distinguish between assigning to access objects and assigning to actual arrays. For example:

Listing 62: show array assignments.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Array_Assignments is
3
4
      type Integer_Array is
5
         array (Positive range <>) of Integer;
6
7
      type Integer_Array_Access is
8
         access Integer_Array;
9
10
      procedure Show_Array
11
         (Name : String;
12
         Arr : Integer_Array_Access) is
13
      begin
14
          Put (Name);
15
          for E of Arr.all loop
16
             Put (Integer'Image (E));
17
          end loop;
18
19
          New Line;
20
      end Show_Array;
21
      Arr_1 : constant Integer_Array_Access :=
22
                           new Integer_Array (1 .. 6);
23
```

```
Arr_2 :
                         Integer_Array_Access :=
24
                           new Integer_Array (1 .. 6);
25
   begin
26
      Arr_1.all := (1,
                          2, 3, 5, 8, 13);
27
      Arr_2.all := (21, 34, 55, 89, 144, 233);
28
29
       -- Array assignment
30
      Arr_2.all := Arr_1.all;
31
32
      Show_Array ("Arr_2", Arr_2);
33
34
       -- Access value assignment
35
      Arr_2 := Arr_1;
36
37
      Arr_1.all := (377, 610, 987, 1597, 2584, 4181);
38
39
      Show_Array ("Arr_2", Arr_2);
40
   end Show_Array_Assignments;
41
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Dereferencing.Array_ ⇔Assignments MD5: 9b1f99af081000c28a6bf9b033127ea3

Runtime output

Arr_2 1 2 3 5 8 13 Arr_2 377 610 987 1597 2584 4181

Here, Arr_2.all := Arr_1.all is an array assignment, while Arr_2 := Arr_1 is an access value assignment. By forcing the usage of the .all suffix, the distinction is clear. Implicit dereferencing, however, could be confusing here. (For example, the .all suffix in Arr_2 := Arr_1.all is an oversight by the programmer when the intention actually was to use access values on both sides.) Therefore, implicit dereferencing is only supported in those cases where there's no risk of ambiguities or oversights.

Records

Let's see an example of implicit dereferencing of a record:

```
Listing 63: show_dereferencing.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_Dereferencing is
3
4
       type Rec is record
5
          I : Integer;
6
          F : Float;
7
      end record;
8
9
      type Rec Access is access Rec;
10
11
      R : constant Rec Access := new Rec;
12
   begin
13
      R.all := (I => 1, F => 5.0);
14
15
      Put Line ("R.I: "
16
                 & Integer'Image (R.I));
17
```

```
18 Put_Line ("R.F: "
19 & Float'Image (R.F));
20 end Show_Dereferencing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Dereferencing.

→Record_Implicit_Dereferencing

MD5: 9af72502d04f128785f77dcc829d5d48
```

Runtime output

R.I: 1 R.F: 5.00000E+00

Again, we can replace R.all.I by R.I, as record components are implicitly dereferenced. Also, we could use implicit dereference when assigning to record components individually:

R.I := 1; R.F := 5.0;

However, we have to write R.all when assigning to the whole record R.

Attributes

Finally, let's see an example of implicit dereference when using attributes:

Listing 64: show_dereferencing.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Dereferencing is
3
4
      type Integer Array is
5
        array (Positive range <>) of Integer;
6
7
      type Integer_Array_Access is
8
        access Integer_Array;
9
10
      Arr : constant Integer_Array_Access :=
11
                         new Integer_Array (1 .. 6);
12
   begin
13
      Put_Line
14
         ("Arr'First: "
15
         & Integer'Image (Arr'First));
16
      Put Line
17
         ("Arr'Last: "
18
         & Integer'Image (Arr'Last));
19
20
      Put_Line
21
         ("Arr'Component_Size: "
22
         & Integer'Image (Arr'Component_Size));
23
      Put Line
24
         ("Arr.all'Component Size: "
25
         & Integer'Image (Arr.all'Component_Size));
26
27
      Put Line
28
         ("Arr'Size: "
29
         & Integer'Image (Arr'Size));
30
      Put_Line
31
```

```
32 ("Arr.all'Size: "
33 & Integer'Image (Arr.all'Size));
34 end Show_Dereferencing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Dereferencing.Array_

□Implicit_Dereferencing

MD5: 5730e18c8d2ed5e26a4d7d325a46a7e9
```

Runtime output

```
Arr'First: 1
Arr'Last: 6
Arr'Component_Size: 32
Arr.all'Component_Size: 32
Arr'Size: 128
Arr.all'Size: 192
```

Here, we can write Arr'First and Arr'Last instead of Arr.all'First and Arr. all'Last, respectively, because Arr is implicitly dereferenced. The same applies to Arr'Component_Size. Note that we can write both Arr'Size and Arr.all'Size, but they have different meanings:

- Arr'Size is the size of the access object; while
- Arr.all'Size indicates the size of the actual array Arr.

In other words, the Size attribute is *not* implicitly dereferenced. In fact, any attribute that could potentially be ambiguous is not implicitly dereferenced. Therefore, in those cases, we must explicitly indicate (by using .all or not) how we want to use the attribute.

Summary

The following table summarizes all instances where implicit dereferencing is supported:

Entities	Standard Usage	Implicit Dereference	
Array components	Arr.all (I)	Arr (I)	
Array slices	Arr.all (F L)	Arr (F L)	
Record components	Rec.all.C	Rec.C	
Array attributes	Arr.all'First	Arr'First	
	Arr.all'First (N)	Arr'First (N)	
	Arr.all'Last	Arr'Last	
	Arr.all'Last (N)	Arr'Last (N)	
	Arr.all'Range	Arr'Range	
	Arr.all'Range (N)	Arr'Range (N)	
	Arr.all'Length	Arr'Length	
	Arr.all'Length (N)	Arr'Length (N)	
	Arr.all'Component_Size	Arr'Component_Size	
Task attributes	T.all'Identity	T'Identity	
	T.all'Storage_Size	T'Storage_Size	
	T.all'Terminated	T'Terminated	
	T.all'Callable	T'Callable	
Tagged type attributes	X.all'Tag	X'Tag	
Other attributes	X.all'Valid	X'Valid	
	X.all'Old X'Old		
	A.all'Constrained	A'Constrained	

1 In the Ada Reference Manual

- 4.1 Names²⁶⁵
- 4.1.1 Indexed Components²⁶⁶
- 4.1.2 Slices²⁶⁷
- 4.1.3 Selected Components²⁶⁸
- 4.1.4 Attributes²⁶⁹

15.8 Ragged arrays

Ragged arrays — also known as jagged arrays — are non-uniform, multidimensional arrays. They can be useful to implement tables with varying number of coefficients, as we discuss as an example in this section.

15.8.1 Uniform multidimensional arrays

Consider an algorithm that processes data based on coefficients that depends on a selected quality level:

Quality level	Number of coefficients	#1	#2	#3	#4	#5
Simplified	1	0.15				
Better	3	0.02	0.16	0.27		
Best	5	0.01	0.08	0.12	0.20	0.34

(Note that this is just a bogus table with no real purpose, as we're not trying to implement any actual algorithm.)

We can implement this table as a two-dimensional array (Calc_Table), where each quality level has an associated array:

Listing 65: data_processing.ads

```
package Data Processing is
1
2
      type Quality Level is
3
         (Simplified, Better, Best);
4
5
   private
6
7
      Calc_Table : constant array
8
         (Quality_Level, 1 .. 5) of Float :=
9
           (Simplified =>
10
                (0.15, 0.00, 0.00, 0.00, 0.00),
11
            Better
                        =>
12
                (0.02, 0.16, 0.27, 0.00, 0.00),
13
            Best
                        =>
14
                (0.01, 0.08, 0.12, 0.20, 0.34));
15
16
```

(continues on next page)

²⁶⁵ http://www.ada-auth.org/standards/22rm/html/RM-4-1.html

²⁶⁶ http://www.ada-auth.org/standards/22rm/html/RM-4-1-1.html

²⁶⁷ http://www.ada-auth.org/standards/22rm/html/RM-4-1-2.html

²⁶⁸ http://www.ada-auth.org/standards/22rm/html/RM-4-1-3.html

²⁶⁹ http://www.ada-auth.org/standards/22rm/html/RM-4-1-4.html

```
17 Last : constant array
18 (Quality_Level) of Positive :=
19 (Simplified => 1,
20 Better => 3,
21 Best => 5);
22
23 end Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Ragged_Arrays.

⊸Uniform_Table

MD5: befa8d2b684ee20495f2dd6907dc44d4
```

Note that, in this implementation, we have a separate table Last that indicates the actual number of coefficients of each quality level.

Alternatively, we could use a record (Table_Coefficient) that stores the number of coefficients and the actual coefficients:

Listing 66: data_processing.ads

```
package Data_Processing is
1
2
      type Quality_Level is
3
         (Simplified, Better, Best);
4
5
      type Data is
6
         array (Positive range <>) of Float;
7
8
   private
9
10
      type Table_Coefficient is record
11
         Last : Positive;
12
          Coef : Data (1 .. 5);
13
      end record;
14
15
      Calc Table : constant array
16
         (Quality_Level) of Table_Coefficient :=
17
           (Simplified =>
18
                (1, (0.15, 0.00, 0.00, 0.00, 0.00)),
19
            Better
                        =>
20
                (3, (0.02, 0.16, 0.27, 0.00, 0.00)),
21
            Best
22
                       =>
                (5, (0.01, 0.08, 0.12, 0.20, 0.34)));
23
24
   end Data_Processing;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Ragged_Arrays.

⊲Uniform_Table

MD5: 4c8602f6ecede0ac1231838c0a0a54b7
```

In this case, we have a unidimensional array where each component (of Table_Coefficient type) contains an array (Coef) with the coefficients.

This is an example of a Process procedure that references the Calc_Table:

Listing 67: data_processing-operations.ads

```
package Data_Processing.Operations is
procedure Process (D : in out Data;
Q : Quality_Level);
end Data Processing.Operations;
```

Listing 68: data processing-operations.adb

```
package body Data Processing.Operations is
1
2
      procedure Process (D : in out Data;
3
                           Q :
                                      Quality_Level) is
4
      begin
5
          for I in D'Range loop
6
             for J in 1 .. Calc_Table (Q).Last loop
7
               -- ... * Calc_Table (Q).Coef (J)
8
               null:
9
             end loop;
10
             -- D (I) := ...
11
             null;
12
         end loop;
13
      end Process;
14
15
   end Data_Processing.Operations;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Ragged_Arrays.

→Uniform_Table

MD5: 2b0d2cee265509e64e507cfa6289bdcc
```

Note that, to loop over the coefficients, we're using **for** J **in** 1 ... Calc_Table (Q). Last **loop** instead of **for** J **in** Calc_Table (Q) 'Range **loop**. As we're trying to make a non-uniform array fit in a uniform array, we cannot simply loop over all elements using the **Range** attribute, but must be careful to use the correct number of elements in the loop instead.

Also, note that Calc_Table has 15 coefficients in total. Out of those coefficients, 6 coefficients (or 40 percent of the table) aren't being used. Naturally, this is wasted memory space. We can improve this by using ragged arrays.

15.8.2 Non-uniform multidimensional array

Ragged arrays are declared by using an access type to an array. By doing that, each array can be declared with a different size, thereby creating a non-uniform multidimensional array.

For example, we can declare a constant array Table as a ragged array:

Listing 69: data_processing.ads

```
type Integer_Array_Access is
8
         access constant Integer_Array;
9
10
       Table : constant array (1 .. 3) of
11
                  Integer_Array_Access :=
12
         (1 \Rightarrow new Integer_Array'(1 \Rightarrow 15),
13
          2 => new Integer_Array'(1 => 12,
14
                                      2 => 15,
15
                                      3 => 20)
16
          3 => new Integer_Array'(1 => 12,
17
                                       2 => 15,
18
                                      3 => 20,
19
                                      4 => 20,
20
                                      5 => 25,
21
                                      6 => 30));
22
23
   end Data_Processing;
24
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Ragged_Arrays.

⊸Simple_Ragged_Array

MD5: 28e044a43bf45585a0268c60d63c629e
```

Here, each component of Table is an access to another array. As each array is allocated via **new**, those arrays may have different sizes.

We can rewrite the example from the previous subsection using a ragged array for the Calc_Table:

Listing 70: data_processing.ads

```
package Data_Processing is
1
2
      type Quality_Level is
3
         (Simplified, Better, Best);
4
5
       type Data is
6
         array (Positive range <>) of Float;
7
8
   private
9
10
       type Coefficients is access constant Data;
11
12
      Calc_Table : constant array (Quality_Level) of
13
                       Coefficients :=
14
         (Simplified =>
15
              new Data'(1 => 0.15),
16
          Better
                      =>
17
              new Data'(0.02, 0.16, 0.27),
18
          Best
                      =>
19
              new Data'(0.01, 0.08, 0.12,
20
                         0.20, 0.34));
21
22
   end Data_Processing;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Ragged_Arrays.

⇔Ragged_Table

MD5: 0781b27cba27dbd1e74da54e425a1f4b
```

Now, we aren't wasting memory space because each data component has the right size that is required for each quality level. Also, we don't need to store the number of coefficients, as this information is automatically available from the array initialization — via the allocation of the Data array for the Coefficients type.

Note that the Coefficients type is defined as **access constant**. We discuss *access-to-constant types* (page 637) in more details later on.

This is the adapted Process procedure:

Listing 71: data_processing-operations.ads

```
package Data_Processing.Operations is
procedure Process (D : in out Data;
Q : Quality_Level);
end Data Processing.Operations;
```

Listing 72: data_processing-operations.adb

```
package body Data_Processing.Operations is
1
2
      procedure Process (D : in out Data;
3
                                      Quality Level) is
                           Q :
4
      begin
5
          for I in D'Range loop
6
             for J in Calc_Table (Q)'Range loop
7
               -- ... * Calc_Table (Q).Coef (J)
8
               null;
9
             end loop;
10
             -- D (I) := ...
11
             null:
12
         end loop;
13
      end Process;
14
15
   end Data Processing.Operations;
16
```

Now, we can simply loop over the coefficients by writing **for** J **in** Calc_Table (Q) 'Range **loop**, as each element of Calc_Table automatically has the correct range.

15.9 Aliasing

The term aliasing²⁷⁰ refers to objects in memory that we can access using more than a single reference. In Ada, if we allocate an object via **new**, we have a potentially aliased object. We can then have multiple references to this object:

Listing 73: show_aliasing.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Aliasing is
3
      type Integer_Access is access Integer;
4
5
6
      A1, A2 : Integer_Access;
7
   begin
      A1 := new Integer;
8
      A2 := A1;
9
10
```

```
<sup>270</sup> https://en.wikipedia.org/wiki/Aliasing_(computing)
```

```
A1.all := 22;
11
      Put_Line ("A1: " & Integer'Image (A1.all));
12
      Put_Line ("A2: " & Integer'Image (A2.all));
13
14
      A2.all := 24;
15
      Put_Line ("A1: " & Integer'Image (A1.all));
16
      Put_Line ("A2: " & Integer'Image (A2.all));
17
   end Show_Aliasing;
18
```

Code block metadata

Project: Courses.Advanced Ada.Resource Management.Access Types.Aliasing.Aliasing →Via Access MD5: 2fde6073cec9823a1a9d93aec82384e1

Runtime output

22 A1: A2: 22 A1: 24 A2: 24

2

6

7

8

9

In this example, we access the object allocated via **new** by using either A1 or A2, as both refer to the same *aliased* object. In other words, A1 or A2 allow us to access the same object in memory.

Important

Note that aliasing is unrelated to renaming. For example, we could use renaming to write a program that looks similar to the one above:

Listing 74: show_renaming.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   procedure Show_Renaming is
3
      A1 : Integer;
4
      A2 : Integer renames A1;
5
   begin
      A1 := 22;
      Put_Line ("A1: " & Integer'Image (A1));
      Put Line ("A2: " & Integer'Image (A2));
10
      A2 := 24;
11
      Put_Line ("A1: " & Integer'Image (A1));
12
      Put_Line ("A2: " & Integer'Image (A2));
13
   end Show Renaming;
14
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Renaming MD5: 99a47d02000b91f7464dffe994fd8ee6

Runtime output

A1: 22

A2: 22

A1: 24 A2: 24

Here, A1 or A2 are two different names for the same object. However, the object itself isn't aliased.



3.10 Access Types²⁷¹

15.9.1 Aliased objects

As we discussed *previously* (page 595), we use **new** to create aliased objects on the heap. We can also use general access types to access objects that were created on the stack.

By default, objects created on the stack aren't aliased. Therefore, we have to indicate that an object is aliased by using the **aliased** keyword in the object's declaration: Obj : **aliased Integer**;

Let's see an example:

Listing 75: show_aliased_obj.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show Aliased Obj is
3
      type Integer_Access is access all Integer;
4
5
      I_Var : aliased Integer;
6
      A1 : Integer_Access;
7
   begin
8
      A1 := I_Var'Access;
9
10
      A1.all := 22;
11
      Put_Line ("A1: " & Integer'Image (A1.all));
12
   end Show Aliased Obj;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Access_ ⇔Aliased_Obj MD5: 98c8e47d7c2b5df8075918b239a8d476

Runtime output

A1: 22

Here, we declare I_Var as an aliased integer variable and get a reference to it, which we assign to A1. Naturally, we could also have two accesses A1 and A2:

```
Listing 76: show_aliased_obj.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_Aliased_Obj is
3
      type Integer_Access is access all Integer;
4
5
      I_Var : aliased Integer;
6
      A1, A2 : Integer_Access;
7
8
   beain
      A1 := I_Var'Access;
9
      A2 := A1;
10
11
      A1.all := 22;
12
      Put_Line ("A1: " & Integer'Image (A1.all));
13
```

(continues on next page)

²⁷¹ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

```
14 Put_Line ("A2: " & Integer'Image (A2.all));
15
16 A2.all := 24;
17 Put_Line ("A1: " & Integer'Image (A1.all));
18 Put_Line ("A2: " & Integer'Image (A2.all));
19
20 end Show_Aliased_Obj;
```

Code block metadata

Runtime output

A1: 22 A2: 22 A1: 24 A2: 24

In this example, both A1 and A2 refer to the I_Var variable.

Note that these examples make use of these two features:

- 1. The declaration of a general access type (Integer_Access) using access all.
- 2. The retrieval of a reference to I_Var using the Access attribute.

In the next sections, we discuss these features in more details.

1 In the Ada Reference Manual

- 3.3.1 Object Declarations²⁷²
- 3.10 Access Types²⁷³

General access modifiers

Let's now discuss how to declare general access types. In addition to the *standard* (pool-specific) access type declarations, Ada provides two access modifiers:

Туре	Declaration		
	type T_Acc is access all T		
Access-to-constant	<pre>type T_Acc is access constant T</pre>		

Let's look at an example:

Listing 77: integer_access_types.ads

```
package Integer_Access_Types is
type Integer_Access is
access Integer;
type Integer_Access_All is
```

(continues on next page)

²⁷² http://www.ada-auth.org/standards/22rm/html/RM-3-3-1.html
 ²⁷³ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

```
7 access all Integer;
8
9 type Integer_Access_Const is
10 access constant Integer;
11
12 end Integer_Access_Types;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Show_

⇔Access_Modifiers

MD5: 98ccaa703194ae88222ccc5a4400e967
```

As we've seen previously, we can use a type such as Integer_Access to allocate objects dynamically. However, we cannot use this type to refer to declared objects, for example. In this case, we have to use an access-to-variable type such as Integer_Access_All. Also, if we want to access constants — or access objects that we want to treat as constants —, we use a type such as Integer_Access_Const.

Access attribute

To get access to a variable or a constant, we make use of the **Access** attribute. For example, I_Var'Access gives us access to the I_Var object.

Let's look at an example of how to use the integer access types from the previous code snippet:

Listing 78: integer_access_types.ads

```
package Integer_Access_Types is
1
2
      type Integer Access is
3
        access Integer;
4
5
      type Integer_Access_All is
6
        access all Integer;
7
8
      type Integer_Access_Const is
9
        access constant Integer;
10
11
      procedure Show;
12
13
   end Integer_Access_Types;
14
```

```
Listing 79: integer_access_types.adb
```

```
use Ada.Text I0;
   with Ada.Text I0;
1
2
   package body Integer Access Types is
3
4
      I Var : aliased
                                 Integer := 0;
5
      Fact : aliased constant Integer := 42;
6
7
      Dyn Ptr
                   : constant Integer_Access
8
                       := new Integer'(30);
9
      I Var Ptr
                   : constant Integer_Access_All
10
                       := I Var'Access;
11
      I_Var_C_Ptr : constant Integer_Access_Const
12
                       := I Var'Access;
13
      Fact Ptr
                  : constant Integer Access Const
14
```

```
:= Fact'Access;
15
16
      procedure Show is
17
      begin
18
          Put_Line ("Dyn_Ptr:
19
                     & Integer'Image (Dyn_Ptr.all));
20
          Put_Line ("I_Var_Ptr:
21
                     & Integer'Image (I_Var_Ptr.all));
22
          Put_Line ("I_Var_C_Ptr:
23
                     & Integer'Image
24
                         (I_Var_C_Ptr.all));
25
          Put_Line ("Fact_Ptr:
26
                     & Integer'Image (Fact_Ptr.all));
27
       end Show;
28
29
   end Integer_Access_Types;
30
```

Listing 80: show_access_modifiers.adb

```
with Integer_Access_Types;
procedure Show_Access_Modifiers is
begin
Integer_Access_Types.Show;
end Show_Access_Modifiers;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Show_ ⇔Access_Modifiers MD5: c9036f060859207ea14354b26dc8b981

Runtime output

Dyn_Ptr:	30
I_Var_Ptr:	0
I_Var_C_Ptr:	0
Fact_Ptr:	42

In this example, Dyn_Ptr refers to a dynamically allocated object, I_Var_Ptr refers to the I_Var variable, and Fact_Ptr refers to the Fact constant. We get access to the variable and the constant objects by using the **Access** attribute.

Also, we declare I_Var_C_Ptr as an access-to-constant, but we get access to the I_Var variable. This simply means the object I_Var_C_Ptr refers to is treated as a constant. Therefore, we can write I_Var := 22;, but we cannot write I_Var_C_Ptr.all := 22;.

```
1 In the Ada Reference Manual
```

```
    3.10.2 Operations of Access Types<sup>274</sup>
```

Non-aliased objects

As mentioned earlier, by default, declared objects — which are allocated on the stack — aren't aliased. Therefore, we cannot get a reference to those objects. For example:

²⁷⁴ http://www.ada-auth.org/standards/22rm/html/RM-3-10-2.html

Listing 81: show_access_error.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show_Access_Error is
3
      type Integer_Access is access all Integer;
4
      I_Var : Integer;
5
      A1
           : Integer Access;
6
   begin
7
      A1 := I_Var'Access;
8
9
      A1.all := 22;
10
      Put Line ("A1: " & Integer'Image (A1.all));
11
   end Show_Access_Error;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Access_Non_
⇔Aliased_Obj
MD5: 2a9904062eea96ae6dc209493d6f20d4
```

Build output

```
show_access_error.adb:8:10: error: prefix of "Access" attribute must be aliased
gprbuild: *** compilation phase failed
```

In this example, the compiler complains that we cannot get a reference to I_Var because I_Var is not aliased.

Ragged arrays using aliased objects

We can use aliased objects to declare *ragged arrays* (page 630). For example, we can rewrite a previous program using aliased constant objects:

Listing 82: data_processing.ads

```
package Data Processing is
1
2
       type Integer Array is
3
         array (Positive range <>) of Integer;
4
5
   private
6
7
      type Integer_Array_Access is
8
         access constant Integer_Array;
9
10
       Tab 1 : aliased constant Integer Array
11
                  := (1 => 15);
12
       Tab_2 : aliased constant Integer_Array
13
                  := (12, 15, 20);
14
       Tab_3 : aliased constant Integer_Array
15
                  := (12, 15, 20,
16
                      20, 25, 30);
17
18
      Table : constant array (1 .. 3) of
19
                 Integer Array Access :=
20
         (1 => Tab 1'Access,
21
          2 => Tab 2'Access,
22
          3 => Tab_3'Access);
23
24
   end Data_Processing;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Ragged_

⇔Array_Aliased_Objs

MD5: 7e284560c447c02628e34bac982d4ad5
```

Here, instead of allocating the constant arrays dynamically via **new**, we declare three aliased arrays (Tab_1, Tab_2 and Tab_3) and get a reference to them in the declaration of Table.

Aliased access objects

It's interesting to mention that access objects can be aliased themselves. Consider this example where we declare the Integer_Access_Access type to refer to an access object:

```
Listing 83: show_aliased_access_obj.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show_Aliased_Access_Obj is
3
4
      type Integer_Access
                                   is
5
        access all Integer;
6
      type Integer_Access_Access is
7
        access all Integer Access;
8
9
      I Var : aliased Integer;
10
            : aliased Integer Access;
11
      Α
             : Integer_Access_Access;
      R
12
13
   begin
      A := I Var'Access;
14
      B := A'Access;
15
16
      B.all.all := 22;
17
      Put Line ("A: " & Integer'Image (A.all));
18
      Put Line ("B: " & Integer'Image (B.all.all));
19
  end Show Aliased Access Obj;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Aliased_

⇔Access

MD5: 77e9be5e29cfb99aef9409728202ba9d
```

Runtime output

```
A: 22
B: 22
```

After the assignments in this example, B refers to A, which in turn refers to I_Var. Note that this code only compiles because we declare A as an aliased (access) object.

15.9.2 Aliased components

Components of an array or a record can be aliased. This allows us to get access to those components:

Listing 84: show_aliased_components.adb

```
with Ada.Text_I0; use Ada.Text_I0;
procedure Show_Aliased_Components is
```

```
4
       type Integer_Access is access all Integer;
5
6
       type Rec is record
7
          I_Var_1 :
                             Integer;
8
          I_Var_2 : aliased Integer;
9
      end record;
10
11
      type Integer_Array is
12
         array (Positive range <>) of aliased Integer;
13
14
      R
          : Rec := (22, 24);
15
      Arr : Integer_Array (1 .. 3) := (others => 42);
16
         : Integer_Access;
17
      Α
18
   begin
       -- A := R.I_Var_1'Access;
19
                           ^ ERROR: cannot access
       - -
20
                                    non-aliased
       - -
21
                                     component
       - -
22
23
      A := R.I_Var_2'Access;
24
      Put_Line ("A: " & Integer'Image (A.all));
25
26
      A := Arr (2) 'Access;
27
      Put_Line ("A: " & Integer'Image (A.all));
28
   end Show_Aliased_Components;
29
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Aliased_ ⊶Components MD5: 5dfaa248caf8e37a4a3a1e1a24973777

Runtime output

A: 24 A: 42

In this example, we get access to the I_Var_2 component of record R. (Note that trying to access the I_Var_1 component would gives us a compilation error, as this component is not aliased.) Similarly, we get access to the second component of array Arr.

Declaring components with the **aliased** keyword allows us to specify that those are accessible via other paths besides the component name. Therefore, the compiler won't store them in registers. This can be essential when doing low-level programming — for example, when accessing memory-mapped registers. In this case, we want to ensure that the compiler uses the memory address we're specifying (instead of assigning registers for those components).

```
    In the Ada Reference Manual
    3.6 Array Types<sup>275</sup>
```

²⁷⁵ http://www.ada-auth.org/standards/22rm/html/RM-3-6.html

15.9.3 Aliased parameters

In addition to aliased objects and components, we can declare *aliased parameters* (page 472), as we already discussed in an earlier chapter. As we mentioned there, aliased parameters are always passed by reference, independently of the type we're using.

The parameter mode indicates which type we must use for the access type:

Parameter mode	Туре	
aliased in	Access-to-constant	
aliased out	Access-to-variable	
aliased in out	Access-to-variable	

Using aliased parameters in a subprogram allows us to get access to those parameters in the body of that subprogram. Let's see an example:

Listing 85: data_processing.ads

```
package Data_Processing is
procedure Proc (I : aliased in out Integer);
end Data_Processing;
```

Listing 86: data_processing.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Data Processing is
3
4
      procedure Show (I : aliased Integer) is
5
                             ^ equivalent to
6
          - -
                               "aliased in Integer"
          - -
7
8
          type Integer Constant Access is
9
            access constant Integer;
10
11
          A : constant Integer_Constant_Access
12
                := I'Access;
13
      begin
14
          Put_Line ("Value : I "
15
                     & Integer'Image (A.all));
16
      end Show;
17
18
      procedure Set One (I : aliased out Integer) is
19
20
          type Integer Access is access all Integer;
21
22
          procedure Local_Set_One (A : Integer_Access)
23
          is
24
          begin
25
             A.all := 1;
26
          end Local_Set_One;
27
28
       begin
29
          Local Set One (I'Access);
30
      end Set One;
31
32
       procedure Proc (I : aliased in out Integer) is
33
34
```

```
type Integer_Access is access all Integer;
35
36
          procedure Add_One (A : Integer_Access) is
37
          begin
38
              A.all := A.all + 1;
39
          end Add_One;
40
41
       begin
42
          Show (I);
43
          Add_One (I'Access);
44
          Show (I);
45
       end Proc;
46
47
   end Data_Processing;
48
```

Listing 87: show_aliased_param.adb

```
with Data_Processing; use Data_Processing;
procedure Show_Aliased_Param is
I : aliased Integer := 22;
begin
Proc (I);
end Show_Aliased_Param;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Aliasing.Aliased_ ←Rec_Component MD5: 076238603036aa51cafcc013f38bc8f3

Runtime output

Value : I 22 Value : I 23

Here, Proc has an **aliased in out** parameter. In Proc's body, we declare the Integer_Access type as an **access all** type. We use the same approach in body of the Set_One procedure, which has an **aliased out** parameter. Finally, the Show procedure has an **aliased in** parameter. Therefore, we declare the Integer_Constant_Access as an **access constant** type.

Note that parameter aliasing has an influence on how arguments are passed to a subprogram when the parameter is of scalar type. When a scalar parameter is declared as aliased, the corresponding argument is passed by reference. For example, if we had declared **procedure** Show (I : Integer), the argument for I would be passed by value. However, since we're declaring it as **aliased Integer**, it is passed by reference.

In the Ada Reference Manual	
 6.1 Subprogram Declarations²⁷⁶ 	
 6.2 Formal Parameter Modes²⁷⁷ 	
 6.4.1 Parameter Associations²⁷⁸ 	

²⁷⁶ http://www.ada-auth.org/standards/22rm/html/RM-6-1.html

²⁷⁷ http://www.ada-auth.org/standards/22rm/html/RM-6-2.html

²⁷⁸ http://www.ada-auth.org/standards/22rm/html/RM-6-4-1.html

15.10 Accessibility Levels and Rules: An Introduction

This section provides an introduction to accessibility levels and accessibility rules. This topic can be very complicated, and by no means do we intend to cover all the details here. (In fact, discussing all the details about accessibility levels and rules could be a long chapter on its own. If you're interested in them, please refer to the Ada Reference Manual.) In any case, the goal of this section is to present the intention behind the accessibility rules and build intuition on how to best use access types in your code.

1 In the Ada Reference Manual

• 3.10.2 Operations of Access Types²⁷⁹

15.10.1 Lifetime of objects

First, let's talk a bit about lifetime of objects²⁸⁰. We assume you understand the concept, so this section is very short.

In very simple terms, the lifetime of an object indicates when an object still has relevant information. For example, if a variable V gets out of scope, we say that its lifetime has ended. From this moment on, V no longer exists.

For example:

Listing 88:	show	lifetime.	adb
-------------	------	-----------	-----

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   procedure Show_Lifetime is
3
      I_Var_1 : Integer := 22;
4
   begin
5
6
       Inner Block : declare
7
          I Var 2 : Integer := 42;
8
      begin
9
          Put_Line ("I_Var_1: "
10
                     & Integer'Image (I_Var_1));
11
          Put_Line ("I_Var_2: "
12
                    & Integer'Image (I_Var_2));
13
14
          -- I_Var_2 will get out of scope
15
             when the block finishes.
16
      end Inner_Block;
17
18
       -- I_Var_2 is now out of scope...
19
20
       Put_Line ("I_Var_1: "
21
                 & Integer'Image (I_Var_1));
22
      Put_Line ("I_Var_2: "
23
                 & Integer'Image (I_Var_2));
24
25
       -- ERROR: lifetime of I_Var_2 has ended!
26
   end Show_Lifetime;
27
```

Code block metadata

²⁷⁹ http://www.ada-auth.org/standards/22rm/html/RM-3-10-2.html

²⁸⁰ https://en.wikipedia.org/wiki/Variable_(computer_science)#Scope_and_extent

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Accessibility_ →Levels_Rules_Introduction.Lifetime MD5: ebe36f12c832ecfe71399b89801808d4

Build output

```
show_lifetime.adb:24:31: error: "I_Var_2" is undefined
gprbuild: *** compilation phase failed
```

In this example, we declare I_Var_1 in the Show_Lifetime procedure, and I_Var_2 in its Inner_Block.

This example doesn't compile because we're trying to use I_Var_2 after its lifetime has ended. However, if such a code could compile and run, the last call to Put_Line would potentially display garbage to the user. (In fact, the actual behavior would be undefined.)

15.10.2 Accessibility Levels

In basic terms, accessibility levels are a mechanism to assess the lifetime of objects (as we've just discussed). The starting point is the library level: this is the base level, and no level can be deeper than that. We start "moving" to deeper levels when we use a library in a subprogram or call other subprograms for example.

Suppose we have a procedure Proc that makes use of a package Pkg, and there's a block in the Proc procedure:

```
package Pkg is
   -- Library level
end Pkg;
with Pkg; use Pkg;
procedure Proc is
   -- One level deeper than
   -- library level
begin
   declare.
      -- Two levels deeper than
      - -
         library level
  begin
      null;
   end;
end Proc;
```

For this code, we can say that:

- the specification of Pkg is at library level;
- the declarative part of Proc is one level deeper than the library level; and
- the block is two levels deeper than the library level.

(Note that this is still a very simplified overview of accessibility levels. Things start getting more complicated when we use information from Pkg in Proc. Those details will become more clear in the next sections.)

The levels themselves are not visible to the programmer. For example, there's no Access_Level attribute that returns an integer value indicating the level. Also, you cannot

write a user message that displays the level at a certain point. In this sense, accessibility levels are assessed relatively to each other: we can only say that a specific operation is at the same or at a deeper level than another one.

15.10.3 Accessibility Rules

The accessibility rules determine whether a specific use of access types or objects is legal (or not). Actually, accessibility rules exist to prevent *dangling references* (page 652), which we discuss later. Also, they are based on the *accessibility levels* (page 646) we discussed earlier.

Code example

As mentioned earlier, the accessibility level at a specific point isn't visible to the programmer. However, to illustrate which level we have at each point in the following code example, we use a prefix (L0, L1, and L2) to indicate whether we're at the library level (L0) or at a deeper level.

Let's now look at the complete code example:

Listing 89: library level.ads

```
package Library_Level is
1
2
      type L0_Integer_Access is
3
        access all Integer;
4
5
      L0_IA : L0_Integer_Access;
6
7
      L0_Var : aliased Integer;
8
9
   end Library Level;
10
```

Listing 90: show_library_level.adb

```
with Library_Level; use Library_Level;
1
2
   procedure Show Library Level is
3
       type L1 Integer Access is
4
         access all Integer;
5
6
      L0 IA 2 : L0 Integer Access;
7
      L1_IA : L1_Integer_Access;
8
9
      L1_Var : aliased Integer;
10
11
      procedure Test is
12
          type L2_Integer_Access is
13
            access all Integer;
14
15
          L2_IA : L2_Integer_Access;
16
17
          L2 Var : aliased Integer;
18
      begin
19
          L1_IA := L2_Var'Access;
20
21
          - -
          - -
                    ILLEGAL: L2 object to
22
          - -
                              L1 access object
23
24
          L2 IA := L2 Var'Access;
25
26
```

```
- -
                     LEGAL: L2 object to
27
                             L2 access object
28
           - -
       end Test;
29
30
31
    begin
       L0_IA := new Integer'(22);
32
33
       - -
                 LEGAL: L0 object to
        - -
34
                          L0 access object
35
36
       L0_IA_2 := new Integer'(22);
37
                     ~~~/
                         ~~~
38
                    LEGAL: L0 object to
39
        - -
                            L0 access object
40
       - -
41
       L0_IA := L1_Var'Access;
42
43
                 ILLEGAL: L1 object to
       - -
44
                            L0 access object
45
        - -
46
       L0_IA_2 := L1_Var'Access;
47
48
                    ILLEGAL: L1 object to
       - -
49
                              L0 access object
50
       - -
51
       L1_IA := L0_Var'Access;
52
                  ~~~~~
53
       - -
                 LEGAL: L0 object to
54
       - -
                         L1 access object
55
        - -
56
       L1_IA := L1_Var'Access;
57
                  ~~~~~
58
                  LEGAL: L1 object to
        - -
59
                          L1 access object
60
61
       L0_IA := L1_IA;
62
63
       - -
                 ILLEGAL: type mismatch
64
65
       L0_IA := L0_Integer_Access (L1_IA);
66
67
        - -
                 ILLEGAL: cannot convert
        - -
68
       - -
                            L1 access object to
69
                            L0 access object
        - -
70
71
       Test;
72
    end Show_Library_Level;
73
```

Code block metadata

Build output

```
show_library_level.adb:20:16: error: non-local pointer cannot point to local object
show_library_level.adb:42:13: error: non-local pointer cannot point to local object
show_library_level.adb:47:15: error: non-local pointer cannot point to local object
show_library_level.adb:62:13: error: expected type "L0_Integer_Access" defined at_
__library_level.ads:3
```

(continues on next page)

(continued from previous page)

In this example, we declare

- in the Library_Level package: the L0_Integer_Access type, the L0_IA access object, and the L0_Var aliased variable;
- in the Show_Library_Level procedure: the L1_Integer_Access type, the L0_IA_2 and L1_IA access objects, and the L1_Var aliased variable;
- in the nested Test procedure: the L2_Integer_Access type, the L2_IA, and the L2_Var aliased variable.

As mentioned earlier, the Ln prefix indicates the level of each type or object. Here, the n value is zero at library level. We then increment the n value each time we refer to a deeper level.

For instance:

- when we declare the L1_Integer_Access type in the Show_Library_Level procedure, that declaration is one level deeper than the level of the Library_Level package so it has the L1 prefix.
- when we declare the L2_Integer_Access type in the Test procedure, that declaration is one level deeper than the level of the Show_Library_Level procedure so it has the L2 prefix.

Types and Accessibility Levels

It's very important to highlight the fact that:

- types themselves also have an associated level, and
- objects have the same accessibility level as their types.

When we declare the L0_IA_2 object in the code example, its accessibility level is at library level because its type (the L0_Integer_Access type) is at library level. Even though this declaration is in the Show_Library_Level procedure — whose declarative part is one level deeper than the library level —, the object itself has the same accessibility level as its type.

Now that we've discussed the accessibility levels of this code example, let's see how the accessibility rules use those levels.

Operations on Access Types

In very simple terms, the accessibility rules say that:

- operations on access types at the same accessibility level are legal;
- assigning or converting to a deeper level is legal;

Otherwise, operations targeting objects at a *less-deep* level are illegal.

For example, L0_IA := **new Integer**'(22) and L1_IA := L1_Var'Access are legal because we're operating at the same accessibility level. Also, L1_IA := L0_Var'Access is legal because L1_IA is at a deeper level than L0_Var'Access.

However, many operations in the code example are illegal. For instance, L0_IA := $L1_Var'Access$ and L0_IA_2 := $L1_Var'Access$ are illegal because the target objects in the assignment are *less* deep.

Note that the L0_IA := L1_IA assignment is mainly illegal because the access types don't match. (Of course, in addition to that, assigning L1_Var'Access to L0_IA is also illegal in terms of accessibility rules.)

Conversion between Access Types

The same rules apply to the conversion between access types. In the code example, the L0_Integer_Access (L1_IA) conversion is illegal because the resulting object is less deep. That being said, conversions on the same level are fine:

Listing 91: show same level conversion.adb

```
procedure Show Same Level Conversion is
1
      type L1 Integer Access is
2
        access all Integer;
3
4
      type L1_B_Integer_Access is
5
        access all Integer;
6
7
      L1 IA
             : L1 Integer Access;
8
      L1 B IA : L1 B Integer Access;
9
10
      L1_Var : aliased Integer;
11
   begin
12
      L1_IA := L1_Var'Access;
13
14
      L1_B_IA := L1_B_Integer_Access (L1_IA);
15
16
       - -
                  LEGAL: conversion from
17
                          L1 access object to
18
                          L1 access object
19
   end Show Same Level Conversion;
20
```

Code block metadata

Here, we're converting from the L1_Integer_Access type to the L1_B_Integer_Access, which are both at the same level.

15.10.4 Accessibility rules on parameters

Note that the accessibility rules also apply to access values as subprogram parameters. For example, compilation fails for this example:

Listing 92: names.ads

```
package Names is
1
2
      type Name is access all String;
3
4
      type Constant Name is
5
        access constant String;
6
7
      procedure Show (N : Constant_Name);
8
9
   end Names;
10
```

Listing 93: names.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   -- with Ada.Characters.Handling;
3
   -- use Ada.Characters.Handling;
4
5
   package body Names is
6
7
      procedure Show (N : Constant_Name) is
8
      begin
9
             for I in N'Range loop
         - -
10
         - -
               N (I) := To_Lower (N (I));
11
          -- end loop;
12
         Put Line ("Name: " & N.all);
13
      end Show;
14
15
   end Names;
16
```

Listing 94: show names.adb

```
1 with Names; use Names;
2
3 procedure Show_Names is
4 S : aliased String := "John";
5 begin
6 Show (S'Access);
7 end Show Names;
```

Code block metadata

Build output

```
show_names.adb:6:10: error: non-local pointer cannot point to local object
gprbuild: *** compilation phase failed
```

In this case, the S'Access cannot be used as the actual parameter for the N parameter of the Show procedure because it's in a deeper level. If we allocate the string via **new**, however, the code compiles as expected:

```
Listing 95: show_names.adb
```

```
with Names; use Names;
procedure Show_Names is
S : Name := new String'("John");
begin
Show (Constant_Name (S));
end Show_Names;
```

Code block metadata

Runtime output

Name: John

This version of the code works because both object and access object have the same level.

15.10.5 Dangling References

An access value that points to a non-existent object is called a dangling reference. Later on, we'll discuss how dangling references may occur using *unchecked deallocation* (page 660).

Dangling references are created when we have an access value pointing to an object whose lifetime has ended, so it becomes a non-existent object. This could occur, for example, when an access value still points to an object X that has gone out of scope.

As mentioned in the previous section, the accessibility rules of the Ada language ensure that such situations never happen! In fact, whenever possible, the compiler applies those rules to detect potential dangling references at compile time. When this detection isn't possible at compile time, the compiler introduces an *accessibility check* (page 521). If this check fails at runtime, it raises a Program_Error exception — thereby preventing that a dangling reference gets used.

Let's see an example of how dangling references could occur:

Listing 96: show_dangling_reference.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Dangling Reference is
3
4
       type Integer_Access is
5
         access all Integer;
6
7
      I Var 1 : aliased Integer := 22;
8
9
             : Integer_Access;
      Α1
10
   begin
11
      A1 := I_Var_1'Access;
12
      Put_Line ("A1.all:
13
                  & Integer'Image (A1.all));
14
15
      Put_Line ("Inner_Block will start now!");
16
17
      Inner_Block : declare
18
19
          -- I Var 2 only exists in Inner Block
20
21
          I Var 2 : aliased Integer := 42;
22
23
24
          -- A2 only exists in Inner_Block
25
          - -
26
          A2
                   : Integer_Access;
27
       begin
28
          A2 := I_Var_1'Access;
29
          Put_Line ("A2.all: "
30
                     & Integer'Image (A2.all);
31
32
          A1 := I_Var_2'Access;
33
               PROBLEM: A1 and Integer_Access type
          - -
34
                         have longer lifetime than
          - -
35
                         I_Var_2
          - -
36
37
          Put_Line ("A1.all: "
```

(continues on next page)

38

```
& Integer'Image (A1.all));
39
40
          A2 := I Var 2'Access;
41
              PROBLEM: A2 has the same lifetime as
          - -
42
                         I_Var_2, but Integer_Access
          - -
43
                         type has a longer lifetime.
          - -
44
45
          Put_Line ("A2.all: "
46
                     & Integer'Image (A2.all));
47
       end Inner_Block;
48
49
      Put_Line ("Inner_Block has ended!");
50
      Put_Line ("A1.all: "
51
                 & Integer'Image (A1.all));
52
53
   end Show_Dangling_Reference;
54
```

Code block metadata

Build output

```
show_dangling_reference.adb:33:13: error: non-local pointer cannot point to local

object

show_dangling_reference.adb:41:13: error: non-local pointer cannot point to local

object

gprbuild: *** compilation phase failed
```

Here, we declare the access objects A1 and A2 of Integer_Access type, and the I_Var_1 and I_Var_2 objects. Moreover, A1 and I_Var_1 are declared in the scope of the Show_Dangling_Reference procedure, while A2 and I_Var_2 are declared in the Inner_Block.

When we try to compile this code, we get two compilation errors due to violation of accessibility rules. Let's now discuss these accessibility rules in terms of lifetime, and see which problems they are preventing in each case.

- In the A1 := I_Var_2'Access assignment, the main problem is that A1 has a longer lifetime than I_Var_2. After the Inner_Block finishes — when I_Var_2 gets out of scope and its lifetime has ended —, A1 would still be pointing to an object that does not longer exist.
- 2. In the A2 := I_Var_2'Access assignment, however, both A2 and I_Var_2 have the same lifetime. In that sense, the assignment may actually look pretty much OK.
 - However, as mentioned in the previous section, Ada also cares about the lifetime of access types. In fact, since the Integer_Access type is declared outside of the Inner_Block, it has a longer lifetime than A2 and I_Var_2.
 - To be more precise, the accessibility rules detect that A2 is an access object of a type that has a longer lifetime than I_Var_2.

At first glance, this last accessibility rule may seem too strict, as both A2 and I_Var_2 have the same lifetime — so nothing bad could occur when dereferencing A2. However, consider the following change to the code:

```
A2 := I_Var_2'Access;
```

```
A1 := A2;
```

```
    PROBLEM: A1 will still be referring
    to I_Var_2 after the
    Inner_Block, i.e. when the
    lifetime of I_Var_2 has
    ended!
```

Here, we're introducing the A1 := A2 assignment. The problem with this is that I_Var_2's lifetime ends when the Inner_Block finishes, but A1 would continue to refer to an I_Var_2 object that doesn't exist anymore — thereby creating a dangling reference.

Even though we're actually not assigning A2 to A1 in the original code, we could have done it. The accessibility rules ensure that such an error is never introduced into the program.

1 For further reading...

In the original code, we can consider the A2 := I_Var_2' Access assignment to be safe, as we're not using the A1 := A2 assignment there. Since we're confident that no error could ever occur in the Inner_Block due to the assignment to A2, we could replace it with A2 := I_Var_2' Unchecked Access, so that the compiler accepts it. We discuss more about the unchecked access attribute *later in this chapter* (page 654).

Alternatively, we could have solved the compilation issue that we see in the A2 := I_Var_2' Access assignment by declaring another access type locally in the Inner Block:

```
Inner_Block : declare
  type Integer_Local_Access is
    access all Integer;

    I_Var_2 : aliased Integer := 42;
    A2        : Integer_Local_Access;
begin
    A2 := I_Var_2'Access;
    -- This assignment is fine because
    -- the Integer_Local_Access type has
    -- the same lifetime as I_Var_2.
end Inner_Block;
```

With this change, A2 becomes an access object of a type that has the same lifetime as I_Var_2 , so that the assignment doesn't violate the rules anymore.

(Note that in the Inner_Block, we could have simply named the local access type Integer_Access instead of Integer_Local_Access, thereby masking the Integer_Access type of the outer block.)

We discuss the effects of dereferencing dangling references later in this chapter (page 662).

15.11 Unchecked Access

In this section, we discuss the Unchecked_Access attribute, which we can use to circumvent accessibility issues for objects in specific cases. (Note that this attribute only exists for objects, not for subprograms.)

We've seen *previously* (page 645) that the accessibility levels verify the lifetime of access types. Let's see a simplified version of a code example from that section:

Listing 97: integers.ads

```
package Integers is
type Integer_Access is access all Integer;
end Integers;
```

Listing 98: show_access_issue.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Integers; use Integers;
3
4
   procedure Show Access Issue is
5
      I_Var : aliased Integer := 42;
6
7
      Α
             : Integer_Access;
8
   begin
9
      A := I_Var'Access;
10
           PROBLEM: A has the same lifetime as I_Var,
11
       - -
       - -
                     but Integer_Access type has a
12
                     longer lifetime.
13
       - -
14
      Put Line ("A.all: " & Integer'Image (A.all));
15
   end Show_Access_Issue;
16
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_Access. →Dangling_Reference_Rules MD5: 646acabf3f388b52809349463d20d314

Build output

show_access_issue.adb:10:09: error: non-local pointer cannot point to local object
gprbuild: *** compilation phase failed

Here, the compiler complains about the A := $I_Var'Access$ assignment because the Integer_Access type has a longer lifetime than A. However, we know that this assignment to A — and further uses of A in the code — won't cause dangling references to be created. Therefore, we can assume that assigning the access to I_Var to A is safe.

When we're sure that an access assignment cannot possibly generate dangling references, we can the use Unchecked_Access attribute. For instance, we can use this attribute to circumvent the compilation error in the previous code example, since we know that the assignment is actually safe:

```
Listing 99: integers.ads
```

```
package Integers is
type Integer_Access is access all Integer;
end Integers;
```

Listing 100: show_access_issue.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Integers; use Integers;
```

```
4
   procedure Show_Access_Issue is
5
      I_Var : aliased Integer := 42;
6
7
             : Integer_Access;
      Α
8
   begin
9
      A := I_Var'Unchecked_Access;
10
           OK: assignment is now accepted.
11
12
      Put_Line ("A.all: " & Integer'Image (A.all));
13
  end Show Access Issue;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_Access.

→Dangling_Reference_Rules

MD5: a71b9076d9e2983ffb9811183afdf6c1
```

Runtime output

A.all: 42

When we use the Unchecked_Access attribute, most rules still apply. The only difference to the standard **Access** attribute is that unchecked access applies the rules as if the object we're getting access to was being declared at library level. (For the code example we've just seen, the check would be performed as if I_Var was declared in the Integers package instead of being declared in the procedure.)

It is strongly recommended to avoid unchecked access in general. You should only use it when you can safely assume that the access object will be discarded before the object we had access to gets out of scope. Therefore, if this situation isn't clear enough, it's best to avoid unchecked access. (Later in this chapter, we'll see some of the nasty issues that arrive from creating dangling references.) Instead, you should work on improving the software design of your application by considering alternatives such as using containers or encapsulating access types in well-designed abstract data types.

```
1 In the Ada Reference Manual
```

Unchecked Access Value Creation²⁸¹

15.12 Unchecked Deallocation

So far, we've seen multiple examples of using **new** to allocate objects. In this section, we discuss how to manually deallocate objects.

Our starting point to manually deallocate an object is the generic Ada. Unchecked_Deallocation procedure. We first instantiate this procedure for an access type whose objects we want to be able to deallocate. For example, let's instantiate it for the Integer_Access type:

Listing 101: integer_types.ads

```
with Ada.Unchecked_Deallocation;
```

```
1
2
```

```
package Integer_Types is
```

(continues on next page)

²⁸¹ http://www.ada-auth.org/standards/22rm/html/RM-13-10.html

```
4
      type Integer_Access is access Integer;
5
6
7
      -- Instantiation of Ada.Unchecked_Deallocation
8
          for the Integer_Access type:
9
      - -
10
      procedure Free is
11
        new Ada.Unchecked_Deallocation
12
           (Object => Integer,
13
           Name
                 => Integer_Access);
14
  end Integer_Types;
15
```

Code block metadata

Here, we declare the Free procedure, which we can then use to deallocate objects that were allocated for the Integer Access type.

Ada.Unchecked_Deallocation is a generic procedure that we can instantiate for access types. When declaring an instance of Ada.Unchecked_Deallocation, we have to specify arguments for:

- the formal Object parameter, which indicates the type of actual objects that we want to deallocate; and
- the formal Name parameter, which indicates the access type.

In a type declaration such as **type Integer_Access is access Integer**, **Integer** denotes the Object, while Integer_Access denotes the Name.

Because each instance of Ada.Unchecked_Deallocation is bound to a specific access type, we cannot use it for another access type, even if the type we use for the Object parameter is the same:

Listing 102: integer_types.ads

```
with Ada.Unchecked_Deallocation;
1
2
   package Integer_Types is
3
4
      type Integer_Access is access Integer;
5
6
      procedure Free is
7
         new Ada.Unchecked Deallocation
8
           (Object => Integer,
9
                 => Integer Access);
            Name
10
11
      type Another_Integer_Access is access Integer;
12
13
      procedure Free is
14
         new Ada.Unchecked_Deallocation
15
           (Object => Integer,
16
           Name => Another_Integer_Access);
17
   end Integer_Types;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_ →Deallocation.Simple_Unchecked_Deallocation MD5: b9bc58ff60632287237e2e322fcbc63e

Here, we're declaring two Free procedures: one for the Integer_Access type, another for the Another_Integer_Access. We cannot use the Free procedure for the Integer_Access type when deallocating objects associated with the Another_Integer_Access type, even though both types are declared as **access Integer**.

Note that we can use any name when instantiating the Ada.Unchecked_Deallocation procedure. However, naming it Free is very common.

Now, let's see a complete example that includes object allocation and deallocation:

```
Listing 103: integer types.ads
```

```
with Ada.Unchecked_Deallocation;
1
2
   package Integer_Types is
3
4
      type Integer_Access is access Integer;
5
6
      procedure Free is
7
         new Ada.Unchecked_Deallocation
8
           (Object => Integer,
9
            Name
                 => Integer_Access);
10
11
      procedure Show_Is_Null (I : Integer_Access);
12
13
   end Integer_Types;
14
```

Listing 104: integer_types.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Integer Types is
3
4
      procedure Show Is Null (I : Integer Access) is
5
      begin
6
          if I = null then
7
             Put Line ("access value is null.");
8
          else
9
             Put Line ("access value is NOT null.");
10
          end if;
11
      end Show Is Null;
12
13
   end Integer_Types;
14
```

Listing 105: show_unchecked_deallocation.adb

```
with Ada.Text IO;
                         use Ada.Text_I0;
1
   with Integer_Types; use Integer_Types;
2
3
   procedure Show_Unchecked_Deallocation is
4
5
      I : Integer_Access;
6
7
   begin
8
      Put ("We haven't called new yet... ");
9
      Show_Is_Null (I);
10
11
```

```
Put ("Calling new...");
I I := new Integer;
Show_Is_Null (I);
Put ("Calling Free...");
Free (I);
Show_Is_Null (I);
end Show_Unchecked_Deallocation;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_ →Deallocation.Unchecked_Deallocation MD5: a9f2df04e2fe0d5ee8c17249b4ae315a

Runtime output

```
We haven't called new yet... access value is null.
Calling new... access value is NOT null.
Calling Free... access value is null.
```

In the Show_Unchecked_Deallocation procedure, we first allocate an object for I and then call Free (I) to deallocate it. Also, we call the Show_Is_Null procedure at three different points: before any allocation takes place, after allocating an object for I, and after deallocating that object.

When we deallocate an object via a call to Free, the corresponding access value — which was previously pointing to an existing object — is set to **null**. Therefore, I = null after the call to Free, which is exactly what we see when running this example code.

Note that it is OK to call Free multiple times for the same access object:

Listing 106: show_unchecked_deallocation.adb

```
with Integer_Types; use Integer_Types;
1
2
   procedure Show_Unchecked_Deallocation is
3
4
      I : Integer_Access;
5
6
   begin
7
      I := new Integer;
8
9
      Free (I);
10
       Free (I);
11
       Free (I);
12
   end Show_Unchecked_Deallocation;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_

→Deallocation.Unchecked_Deallocation

MD5: ce7f4f912f12d723ca673ca36a478765
```

The multiple calls to Free for the same access object don't cause any issues. Because the access value is null after the first call to Free (I), we're actually just passing **null** as an argument in the second and third calls to Free. However, any attempt to deallocate an access value of null is ignored in the Free procedure, so the second and third calls to Free don't have any effect.

1 In the Ada Reference Manual

- 4.8 Allocators²⁸²
- 13.11.2 Unchecked Storage Deallocation²⁸³

15.12.1 Unchecked Deallocation and Dangling References

We've discussed *dangling references* (page 652) before. In this section, we discuss how unchecked deallocation can create dangling references and the issues of having them in an application.

Let's reuse the last example and introduce I_2, which will point to the same object as I:

Listing 107: show_unchecked_deallocation.adb

```
with Integer_Types; use Integer_Types;
1
2
   procedure Show Unchecked Deallocation is
3
4
      I, I 2 : Integer Access;
5
6
   begin
7
8
      I := new Integer;
9
      I 2 := I;
10
11
       -- NOTE: I_2 points to the same
12
                 object as I.
       - -
13
14
15
       -- Use I and I 2...
16
17
          ... then deallocate memory...
18
19
20
      Free (I);
21
22
       -- NOTE: at this point, I_2 is a
23
                  dangling reference!
24
       - - -
25
       -- Further calls to Free (I)
26
       -- are OK!
27
28
      Free (I);
29
      Free (I);
30
31
       -- A call to Free (I 2) is
32
       -- NOT OK:
33
34
      Free (I_2);
35
   end Show_Unchecked_Deallocation;
36
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_ →Deallocation.Unchecked_Deallocation MD5: ee5c20209a113a6c1bc7895b8ebdb174

²⁸² http://www.ada-auth.org/standards/22rm/html/RM-4-8.html

²⁸³ http://www.ada-auth.org/standards/22rm/html/RM-13-11-2.html

Runtime output

free(): double free detected in tcache 2

raised PROGRAM_ERROR : unhandled signal

As we've seen before, we can have multiple calls to Free (I). However, the call to Free (I_2) is bad because I_2 is not null. In fact, it is a dangling reference — i.e. I_2 points to an object that doesn't exist anymore. Also, the first call to Free (I) will reclaim the storage that was allocated for the object that I originally referred to. The call to Free (I_2) will then try to reclaim the previously-reclaimed object, but it'll fail in an undefined manner.

Because of these potential errors, you should be very careful when using unchecked deallocation: it is the programmer's responsibility to avoid creating dangling references!

For the example we've just seen, we could avoid creating a dangling reference by explicitly assigning **null** to I_2 to indicate that it doesn't point to any specific object:

Listing 108: show_unchecked_deallocation.adb

```
with Integer_Types; use Integer_Types;
1
2
   procedure Show Unchecked Deallocation is
3
4
       I, I_2 : Integer_Access;
5
6
   begin
7
       I := new Integer;
8
9
       I 2 := I;
10
11
       -- NOTE: I 2 points to the same
12
                  object as I.
13
       - -
14
15
       -- Use I and I_2...
16
17
       - -
           ... then deallocate memory...
       - -
18
19
20
       I 2 := null;
21
22
       -- NOTE: now, I 2 doesn't point to
23
                  any object, so calling
       - -
24
       - -
                  Free (I 2) is OK.
25
26
       Free (I);
27
       Free (I_2);
28
   end Show_Unchecked_Deallocation;
29
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_

→Deallocation.Unchecked_Deallocation

MD5: 3381ba594cbbc0f1547e3f819bae0f97
```

Now, calling Free (I_2) doesn't cause any issues because it doesn't point to any object.

Note, however, that this code example is just meant to illustrate the issues of dangling pointers and how we could circumvent them. We're not suggesting to use this approach when designing an implementation. In fact, it's not practical for the programmer to make every possible dangling reference become null if the calls to Free are strewn throughout the code.

The suggested design is to not use Free in the client code, but instead hide its use within bigger abstractions. In that way, all the occurrences of the calls to Free are in one package, and the programmer of that package can then prevent dangling references. We'll discuss these *design strategies* (page 669) later on.

15.12.2 Dereferencing dangling references

Of course, you shouldn't try to dereference a dangling reference because your program becomes erroneous, as we discuss in this section. Let's see an example:

Listing 109: show_unchecked_deallocation.adb

```
with Ada.Text_I0;
                         use Ada.Text I0;
1
   with Integer_Types; use Integer_Types;
2
3
   procedure Show Unchecked Deallocation is
4
5
      I_1, I_2 : Integer_Access;
6
7
   begin
8
      I 1 := new Integer'(42);
9
      I_2 := I_1;
10
11
      Put_Line ("I 1.all = "
12
                 & Integer'Image (I_1.all));
13
       Put_Line ("I_2.all = "
14
                 & Integer'Image (I 2.all));
15
16
      Put_Line ("Freeing I_1");
17
      Free (I_1);
18
19
      if I_1 /= null then
20
          Put_Line ("I_1.all = "
21
                     & Integer'Image (I_1.all));
22
      end if;
23
24
       if I 2 /= null then
25
          Put_Line ("I_2.all = "
26
                     & Integer'Image (I 2.all));
27
       end if;
28
   end Show_Unchecked_Deallocation;
29
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_ →Deallocation.Unchecked_Deallocation MD5: 8536190aa5bbafa715ad8153aaeb4889

Runtime output

I_1.all = 42 I_2.all = 42 Freeing I_1 I_2.all = 1743404846

In this example, we allocate an object for I_1 and make I_2 point to the same object. Then, we call Free (I), which has the following consequences:

- The call to Free (I_1) will try to reclaim the storage for the original object (I_1.all), so it may be reused for other allocations.
- I_1 = **null** after the call to Free (I_1).

- I_2 becomes a dangling reference by the call to Free (I_1).
 - In other words, I_2 is still non-null, and what it points to is now undefined.

In principle, we could check for **null** before trying to dereference the access value. (Remember that when deallocating an object via a call to Free, the corresponding access value is set to **null**.) In fact, this strategy works fine for I_1, but it doesn't work for I_2 because the access value is not **null**. As a consequence, the application tries to dereference I_2.

Dereferencing a dangling reference is erroneous: the behavior is undefined in this case. For the example we've just seen,

- I_2.all might make the application crash;
- I_2.all might give us a different value than before;
- I_2.all might even give us the same value as before (42) if the original object is still available.

Because the effect is unpredictable, it might be really difficult to debug the application and identify the cause.

Having dangling pointers in an application should be avoided at all costs! Again, it is the programmer's responsibility to be very careful when using unchecked deallocation: avoid creating dangling references!

In the Ada Reference Manual

- 13.9.1 Data Validity²⁸⁴
- 13.11.2 Unchecked Storage Deallocation²⁸⁵

15.12.3 Restrictions for Ada.Unchecked_Deallocation

There are two unsurprising restrictions for Ada.Unchecked_Deallocation:

- 1. It cannot be instantiated for access-to-constant types; and
- 2. It cannot be used when the Storage_Size aspect of a type is zero (i.e. when its storage pool is empty).

(Note that this last restriction also applies to the allocation via new.)

Let's see an example of these restrictions:

Listing 110: show_unchecked_deallocation_errors.adb

```
with Ada.Unchecked Deallocation;
1
   procedure Show Unchecked Deallocation Errors is
3
4
      type Integer Access Zero is access Integer
5
        with Storage_Size => 0;
6
7
      procedure Free is
8
        new Ada.Unchecked Deallocation
9
           (Object => Integer,
10
                 => Integer Access Zero);
           Name
11
12
      type Constant Integer Access is
13
        access constant Integer;
14
```

(continues on next page)

²⁸⁴ http://www.ada-auth.org/standards/22rm/html/RM-13-9-1.html

²⁸⁵ http://www.ada-auth.org/standards/22rm/html/RM-13-11-2.html

```
15
       -- ERROR: Cannot use access-to-constant type
16
                  for Name
17
       - -
      procedure Free is
18
        new Ada.Unchecked_Deallocation
19
           (Object => Integer,
20
                  => Constant_Integer_Access);
            Name
21
22
      I : Integer_Access_Zero;
23
24
   begin
25
       - -
          ERROR: Cannot allocate objects from
26
27
                   empty storage pool
      I := new Integer;
28
29
       -- ERROR: Cannot deallocate objects from
30
                   empty storage pool
31
       Free (I);
32
   end Show_Unchecked_Deallocation_Errors;
33
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Unchecked_

→Deallocation.Unchecked_Deallocation_Error

MD5: 5032d13b2eb6b7ca1979282ddd6df98a
```

Build output

Here, we see that trying to instantiate Ada.Unchecked_Deallocation for the Constant_Integer_Access type is rejected by the compiler. Similarly, we cannot allocate or deallocate an object for the Integer_Access_Zero type because its storage pool is empty.

15.13 Null & Not Null Access

\rm 1 Note

This section was originally written by Robert A. Duff and published as Gem #23: Null Considered Harmful²⁸⁶ and Gem #24²⁸⁷.

Ada, like many languages, defines a special **null** value for access types. All values of an access type designate some object of the designated type, except for **null**, which does not designate any object. The null value can be used as a special flag. For example, a singly-linked list can be null-terminated. A Lookup function can return **null** to mean "not found", presuming the result is of an access type:

²⁸⁶ https://www.adacore.com/gems/ada-gem-23

²⁸⁷ https://www.adacore.com/gems/ada-gem-24

Listing 111: show_null_return.ads

```
package Show_Null_Return is
type Ref_Element is access all Element;
Not_Found : constant Ref_Element := null;
function Lookup (T : Table) return Ref_Element;
    -- Returns Not_Found if not found.
end Show_Null_Return;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Null_And_Not_Null_

⇔Access.Null_Return

MD5: 6c4eed750d42685198ec9495805e3e23
```

An alternative design for Lookup would be to raise an exception:

Listing 112: show_not_found_exception.ads

```
package Show_Not_Found_Exception is
Not_Found : exception;
function Lookup (T : Table) return Ref_Element;
.-- Raises Not_Found if not found.
.-- Never returns null.
rend Show_Not_Found_Exception;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Null_And_Not_Null_

⇔Access.Not_Found_Exception

MD5: 6ef47b32d4923838ffc28f43e5db323c
```

Neither design is better in all situations; it depends in part on whether we consider the "not found" situation to be exceptional.

Clearly, the client calling Lookup needs to know whether it can return **null**, and if so, what that means. In general, it's a good idea to document whether things can be null or not, especially for formal parameters and function results. Prior to Ada 2005, we would do that with comments. Since Ada 2005, we can use the **not null** syntax:

```
Listing 113: show_not_null_return.ads
```

```
package Show Not Null Return is
1
      type Ref_Element is access all Element;
2
3
      Not_Found : constant Ref_Element := null;
4
5
      function Lookup (T : Table)
6
                        return not null Ref_Element;
7
      -- Possible since Ada 2005.
8
  end Show_Not_Null_Return;
9
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Null_And_Not_Null_

⇔Access.Not_Null_Return

MD5: 4c0bb95da3b5a7c555a763c4951f7e21
```

This is a complete package for the code snippets above:

```
Listing 114: example.ads
```

```
package Example is
1
2
      type Element is limited private;
3
      type Ref_Element is access all Element;
4
5
      type Table is limited private;
6
7
      Not_Found : constant Ref_Element := null;
8
       function Lookup (T : Table)
9
                         return Ref Element;
10
       -- Returns Not Found if not found.
11
12
      Not_Found_2 : exception;
13
      function Lookup_2 (T : Table)
14
                           return not null Ref_Element;
15
       -- Raises Not_Found_2 if not found.
16
17
      procedure P (X : not null Ref_Element);
18
19
      procedure Q (X : not null Ref_Element);
20
21
   private
22
       type Element is limited
23
          record
24
             Component : Integer;
25
          end record;
26
      type Table is limited null record;
27
   end Example;
28
```

```
Listing 115: example.adb
```

```
package body Example is
1
2
       An_Element : aliased Element;
3
4
       function Lookup (T : Table)
5
                          return Ref_Element is
6
          pragma Unreferenced (T);
7
       begin
8
9
          return Not_Found;
10
       end Lookup;
11
12
       function Lookup_2 (T : Table)
13
                           return not null Ref_Element
14
       is
15
       begin
16
17
          raise Not_Found_2;
18
19
          return An_Element'Access;
20
          -- suppress error: 'missing "return"
21
              statement in function body'
          - -
22
       end Lookup_2;
23
24
       procedure P (X : not null Ref_Element) is
25
       begin
26
          X.all.Component := X.all.Component + 1;
27
```

```
end P;
28
29
       procedure Q (X : not null Ref Element) is
30
       begin
31
           for I in 1 .. 1000 loop
32
              P (X);
33
          end loop;
34
       end Q;
35
36
       procedure R is
37
       begin
38
          Q (An Element'Access);
39
40
       end R;
41
      pragma Unreferenced (R);
42
43
   end Example;
44
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Null_And_Not_Null_

→Access.Complete_Null_Return

MD5: 01895c7d5f843fd215dcc21d807d4187
```

In general, it's better to use the language proper for documentation, when possible, rather than comments, because compile-time and/or run-time checks can help ensure that the "documentation" is actually true. With comments, there's a greater danger that the comment will become false during maintenance, and false documentation is obviously a menace.

In many, perhaps most cases, **null** is just a tripping hazard. It's a good idea to put in **not null** when possible. In fact, a good argument can be made that **not null** should be the default, with extra syntax required when **null** is wanted. This is the way Standard ML^{288} works, for example — you don't get any special null-like value unless you ask for it. Of course, because Ada 2005 needs to be compatible with previous versions of the language, **not null** cannot be the default for Ada.

One word of caution: access objects are default-initialized to **null**, so if you have a **not null** object (or component) you had better initialize it explicitly, or you will get Constraint_Error. **not null** is more often useful on parameters and function results, for this reason.

Another advantage of **not null** over comments is for efficiency. Consider procedures P and Q in this example:

Listing 116:	examp	le-process	ing.ads
--------------	-------	------------	---------

```
package Example.Processing is
procedure P (X : not null Ref_Element);
procedure Q (X : not null Ref_Element);
end Example.Processing;
```

Listing 117: example-processing.adb

```
package body Example.Processing is
```

```
<sup>288</sup> https://en.wikipedia.org/wiki/Standard_ML
```

```
procedure P (X : not null Ref_Element) is
3
      begin
4
          X.all.Component := X.all.Component + 1;
5
      end P;
6
7
      procedure Q (X : not null Ref_Element) is
8
      begin
9
          for I in 1 .. 1000 loop
10
             P (X);
11
          end loop;
12
      end Q;
13
14
   end Example.Processing;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Null_And_Not_Null_

⇔Access.Complete_Null_Return

MD5: dc34b1a27737d57c041be6260dd577fd
```

Without **not null**, the generated code for P will do a check that $X \neq null$, which may be costly on some systems. P is called in a loop, so this check will likely occur many times. With **not null**, the check is pushed to the call site. Pushing checks to the call site is usually beneficial because

- 1. the check might be hoisted out of a loop by the optimizer, or
- 2. the check might be eliminated altogether, as in the example above, where the compiler knows that An_Element 'Access cannot be **null**.

This is analogous to the situation with other run-time checks, such as array bounds checks:

```
Listing 118: show_process_array.ads
```

```
package Show_Process_Array is
1
2
      type My Index is range 1 .. 10;
3
      type My_Array is array (My_Index) of Integer;
4
5
      procedure Process_Array
6
        (X : in out My_Array;
7
         Index :
                        My_Index);
8
9
   end Show_Process_Array;
10
```

```
Listing 119: show_process_array.adb
```

```
package body Show_Process_Array is
1
2
      procedure Process_Array
3
        (X : in out My_Array;
4
         Index :
                       My_Index) is
5
      begin
6
         X (Index) := X (Index) + 1;
7
      end Process_Array;
8
  end Show_Process_Array;
10
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Null_And_Not_Null_ ⇔Access.Process_Array MD5: 32424432f5b2e3013292680f92a04320

If X (Index) occurs inside $Process_Array$, there is no need to check that Index is in range, because the check is pushed to the caller.

15.14 Design strategies for access types

Previously, we learned about *dangling references* (page 652) and discussed the effects of *dereferencing them* (page 662). Also, we've seen the relationship between *unchecked deallocation and dangling references* (page 660). Ensuring that all calls to Free for a specific access type will never cause dangling references can become an arduous task — if not impossible — if those calls are located in different parts of the source code.

Although we used access types directly in the main application in many of the previous code examples from this chapter, this approach was in fact selected just for illustration purposes — i.e. to make the code look simpler. In general, however, we should avoid this approach. Instead, our recommendation is to encapsulate the access types in some form of abstraction. In this section, we discuss design strategies for access types that take this recommendation into account.

15.14.1 Abstract data type for access types

The simplest form of abstraction is of course an abstract data type. For example, we could declare a limited private type, which allows us to hide the access type and to avoid copies of references that could potentially become dangling references. (We discuss limited private types later *in another chapter* (page 787).)

Let's see an example:

```
Listing 120: access_type_abstraction.ads
```

```
package Access_Type_Abstraction is
1
2
      type Info is limited private;
3
4
       function To Info (S : String) return Info;
5
6
       function To String (Obj : Info)
7
                            return String;
8
9
      function Copy (Obj : Info) return Info;
10
11
      procedure Copy (To : in out Info;
12
                        From :
                                      Info):
13
14
      procedure Append (Obj : in out Info;
15
                              : String);
                          S
16
17
      procedure Reset (Obj : in out Info);
18
19
      procedure Destroy (Obj : in out Info);
20
21
   private
22
23
      type Info is access String;
24
25
   end Access Type Abstraction;
26
```

Listing 121: access_type_abstraction.adb

```
with Ada.Unchecked_Deallocation;
1
2
   package body Access_Type_Abstraction is
3
4
       function To Info (S : String) return Info is
5
         (new String'(S));
6
7
       function To_String (Obj : Info)
8
                             return String is
9
         (if Obj /= null then Obj.all else "");
10
11
       function Copy (Obj : Info) return Info is
12
         (To_Info (To_String (Obj)));
13
14
      procedure Copy (To : in out Info;
15
                                       Info) is
                        From :
16
      begin
17
          Destroy (To);
18
          To := Copy (From);
19
20
      end Copy;
21
      procedure Append (Obj : in out Info;
22
                          S
                             : String) is
23
          New_Info : constant Info :=
24
                        To_Info (To_String (Obj) & S);
25
      begin
26
          Destroy (Obj);
27
          Obj := New_Info;
28
      end Append;
29
30
      procedure Reset (Obj : in out Info) is
31
      begin
32
          Destroy (Obj);
33
      end Reset;
34
35
      procedure Destroy (Obj : in out Info) is
36
          procedure Free is
37
            new Ada.Unchecked_Deallocation
38
39
              (Object => String,
                      => Info);
40
               Name
      begin
41
          Free (Obj);
42
      end Destroy;
43
44
   end Access_Type_Abstraction;
45
```

Listing 122: main.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Access_Type_Abstraction;
3
   use Access_Type_Abstraction;
4
5
   procedure Main is
6
      Obj_1 : Info := To_Info ("hello");
7
      Obj_2 : Info := Copy (Obj_1);
8
   begin
9
      Put_Line ("TO_INFO / COPY");
10
      Put_Line ("0bj_1 : "
11
```

```
(continued from previous page)
```

```
& To_String (Obj_1));
12
      Put_Line ("Obj_2 : "
13
                 & To_String (Obj_2));
14
      Put_Line ("----");
15
16
      Reset (Obj_1);
17
      Append (Obj_2, " world");
18
19
      Put_Line ("RESET / APPEND");
20
      Put_Line ("Obj_1 : "
21
                 & To_String (Obj_1));
22
      Put_Line ("Obj_2 :
23
                 & To_String (Obj_2));
24
      Put_Line ("----");
25
26
      Copy (From => Obj_2,
27
                 => Obj_1);
             То
28
29
      Put Line ("COPY");
30
      Put_Line ("Obj_1 : "
31
                 & To_String (Obj_1));
32
       Put_Line ("Obj_2 :
33
                 & To_String (Obj_2));
34
      Put_Line ("-----");
35
36
37
      Destroy (Obj_1);
38
      Destroy (Obj_2);
39
      Put_Line ("DESTROY");
40
      Put_Line ("Obj_1 :
41
                 & To_String (Obj_1));
42
       Put_Line ("Obj_2 : "
43
                 & To_String (Obj_2));
44
      Put_Line ("-----");
45
46
      Append (Obj_1, "hey");
47
48
      Put_Line ("APPEND");
49
      Put_Line ("Obj_1 : "
50
                 & To_String (Obj_1));
51
      Put_Line ("-----");
52
53
      Put_Line ("APPEND");
54
      Append (Obj_1, " there");
55
      Put Line ("Obj 1 : "
56
                 & To_String (Obj_1));
57
58
      Destroy (Obj_1);
59
      Destroy (Obj_2);
60
   end Main;
61
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Design_Strategies.

⇔Access_Type_Abstraction

MD5: a335caeba4f1fb952a2e0d8d6bc52f75
```

Runtime output

TO_INFO / COPY Obj_1 : hello

Obj_2 : hello RESET / APPEND Obj_1 : Obj_2 : hello world -----COPY Obj_1 : hello world Obj_2 : hello world DESTROY Obj_1 : Obj_2 : - - - - - - - - - -APPEND Obj_1 : hey - - - - - - - - - -APPEND Obj_1 : hey there

In this example, we hide an access type in the Info type — a limited private type. We allocate an object of this type in the To_Info function and deallocate it in the Destroy procedure. Also, we make sure that the reference isn't copied in the Copy function — we only copy the designated value in this function. This strategy eliminates the possibility of dangling references, as each reference is encapsulated in an object of Info type.

15.14.2 Controlled type for access types

In the previous code example, the Destroy procedure had to be called to deallocate the hidden access object. We could make sure that this deallocation happens automatically by using a controlled (or limited controlled) type. (We discuss *controlled types* (page 835) in another chapter.)

Let's adapt the previous example and declare Info as a limited controlled type:

Listing 123: access_type_abstraction.ads

```
with Ada.Finalization;
1
2
   package Access_Type_Abstraction is
3
4
      type Info is limited private;
5
6
      function To Info (S : String) return Info;
7
8
      function To_String (Obj : Info)
9
                            return String;
10
11
      function Copy (Obj : Info) return Info;
12
13
      procedure Copy (To : in out Info;
14
                        From : Info);
15
16
      procedure Append (Obj : in out Info;
17
                          S
                                       String);
18
                              :
19
      procedure Reset (Obj : in out Info);
20
21
   private
22
23
```

```
type String_Access is access String;
24
25
      type Info is new
26
         Ada.Finalization.Limited_Controlled with
27
          record
28
             Str_A : String_Access;
29
          end record;
30
31
      procedure Initialize (Obj : in out Info);
32
      procedure Finalize (Obj : in out Info);
33
34
   end Access_Type_Abstraction;
35
```

Listing 124: access_type_abstraction.adb

```
with Ada. Unchecked Deallocation;
1
2
   package body Access Type Abstraction is
3
4
5
          STRING_ACCESS SUBPROGRAMS
6
       - -
7
8
       function To_String_Access (S : String)
9
                                    return String Access
10
       is
11
         (new String'(S));
12
13
       function To_String (S : String_Access)
14
                             return String is
15
         (if S /= null then S.all else "");
16
17
       procedure Free is
18
         new Ada.Unchecked_Deallocation
19
           (Object => String,
20
                  => String_Access);
            Name
21
22
23
       -- PRIVATE SUBPROGRAMS
24
25
26
       procedure Initialize (Obj : in out Info) is
27
       begin
28
              Put_Line ("Initializing Info");
29
          - -
          Obj.Str_A := null;
30
31
          - -
              NOTE: This line has just been added to
          - -
32
                     illustrate the "automatic" call to
          - -
33
          - -
                     Initialize. Actually, this
34
                     assignment isn't needed, as
          - -
35
                     the Str_A component is
          - -
36
                     automatically initialized to null
37
          - -
                     upon object construction.
38
          - -
       end Initialize;
39
40
       procedure Finalize (Obj : in out Info) is
41
       begin
42
             Put Line ("Finalizing Info");
43
          Free (Obj.Str A);
44
       end Finalize;
45
46
```

```
(continued from previous page)
```

```
47
          PUBLIC SUBPROGRAMS
48
       - -
49
       - -
50
       function To_Info (S : String) return Info is
51
         (Ada.Finalization.Limited_Controlled
52
          with Str_A => To_String_Access (S));
53
54
       function To_String (Obj : Info)
55
                             return String is
56
         (To_String (Obj.Str_A));
57
58
       function Copy (Obj : Info) return Info is
59
         (To_Info (To_String (Obj.Str_A)));
60
61
      procedure Copy (To : in out Info;
62
                        From :
                                      Info) is
63
      begin
64
          Free (To.Str_A);
65
          To.Str_A := To_String_Access
66
                         (To_String (From.Str_A));
67
      end Copy;
68
69
      procedure Append (Obj : in out Info;
70
                          S
                             :
                                        String) is
71
          New_Str_A : constant String_Access :=
72
73
                         To_String_Access
                            (To_String (Obj.Str_A) & S);
74
      begin
75
          Free (Obj.Str_A);
76
          Obj.Str_A := New_Str_A;
77
      end Append;
78
79
      procedure Reset (Obj : in out Info) is
80
      begin
81
          Free (Obj.Str_A);
82
      end Reset;
83
84
   end Access_Type_Abstraction;
85
```

Listing 125: main.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_Type_Abstraction;
3
   use Access_Type_Abstraction;
4
5
   procedure Main is
6
      Obj 1 : Info := To Info ("hello");
7
      Obj_2 : Info := Copy (Obj_1);
8
   begin
9
10
       -- TO_INFO / COPY
11
12
      Put_Line ("T0_INF0 / COPY");
13
14
      Put_Line ("Obj_1 : "
15
                 & To_String (Obj_1));
16
      Put Line ("Obj 2 : "
17
                 & To_String (Obj_2));
18
      Put_Line ("-----");
19
```

```
- -
-- RESET: Obj_1
-- APPEND: Obj_2
- -
Put_Line ("RESET / APPEND");
Reset (Obj_1);
Append (Obj_2, " world");
Put_Line ("Obj_1 : "
         & To_String (Obj_1));
Put_Line ("Obj_2 :
         & To_String (Obj_2));
Put_Line ("-----");
-- COPY: Obj_2 => Obj_1
Put_Line ("COPY");
Copy (From => Obj_2,
     To => Obj_1);
Put_Line ("Obj_1 : "
         & To_String (Obj_1));
Put_Line ("Obj_2 : "
         & To_String (Obj_2));
Put_Line ("-----");
-- RESET: Obj_1, Obj_2
Put_Line ("RESET");
Reset (Obj_1);
Reset (Obj_2);
Put_Line ("Obj_1 : "
         & To_String (Obj_1));
Put_Line ("Obj_2 : "
         & To_String (Obj_2));
Put_Line ("-----");
-- COPY: Obj_2 => Obj_1
Put_Line ("COPY");
Copy (From => Obj_2,
         => Obj_1);
     То
Put Line ("Obj 1 : "
         & To_String (Obj_1));
Put Line ("Obj 2 : '
         & To_String (Obj_2));
Put_Line ("-----");
-- APPEND: Obj 1 with "hey"
```

20 21

22

23 24

25 26

27

28 29

30

31

32

33

34 35 36

37 38

39 40

41

42 43

44 45

46

47

48 49 50

51 52

53 54

55

56 57

58

59

60

61

62 63 64

65 66

67 68

69

70 71

72

73

74

75

76 77 78

79 80

```
Put_Line ("APPEND");
81
82
      Append (Obj_1, "hey");
83
84
      Put_Line ("Obj_1 : "
85
                 & To_String (Obj_1));
86
      Put_Line ("----");
87
88
89
       -- APPEND: Obj_1 with "there"
90
91
      Put_Line ("APPEND");
92
93
      Append (Obj_1, " there");
94
95
      Put_Line ("Obj_1 : "
96
                 & To_String (Obj_1));
97
   end Main;
98
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Design_Strategies. ⇔Access_Type_Limited_Controlled_Abstraction MD5: e98659ad1b87be56fb173fa407ab7e82

Runtime output

```
TO INFO / COPY
Obj 1 : hello
Obj_2 : hello
- - - - - - - -
RESET / APPEND
Obj_1 :
Obj_2 : hello world
COPY
Obj 1 : hello world
Obj_2 : hello world
- - - - - - - - - .
RESET
Obj_1 :
Obj_2 :
- - - - - - - -
COPY
Obj_1 :
0bj_2 :
APPEND
Obj 1 : hey
-----
APPEND
Obj_1 : hey there
```

Of course, because we're using the Limited_Controlled type from the Ada. Finalization package, we had to adapt the prototype of the subprograms from the Access_Type_Abstraction. In this version of the code, we only have the allocation taking place in the To_Info procedure, but we don't have a Destroy procedure for deallocation: this call was moved to the Finalize procedure.

Since objects of the Info type — such as Obj_1 in the Show_Access_Type_Abstraction procedure — are now controlled, the Finalize procedure is automatically called when they go out of scope. In this procedure, which we override for the Info type, we perform the deallocation of the internal access object Str_A. (You may uncomment the calls to Put_Line in the body of the Initialize and Finalize subprograms to confirm that these subprograms are called in the background.)

15.15 Access to subprograms

So far in this chapter, we focused mainly on access-to-objects. However, we can use access types to subprograms. This is the topic of this section.

15.15.1 Static vs. dynamic calls

In a typical subprogram call, we indicate the subprogram we want to call statically. For example, let's say we've implemented a procedure Proc that calls a procedure P:

```
Listing 126: p.ads
```

```
procedure P (I : in out Integer);
```

Listing 127: p.adb

```
procedure P (I : in out Integer) is
begin
null;
end P;
```

Listing 128: proc.adb

```
with P;
procedure Proc is
I : Integer := 0;
begin
P (I);
end Proc:
```

Code block metadata

The call to P is statically dispatched: every time Proc runs and calls P, that call is always to the same procedure. In other words, we can determine at compilation time which procedure is called.

In contrast, an access to a subprogram allows us to dynamically indicate which subprogram we want to call. For example, if we change Proc in the code above to receive the access to a subprogram P as a parameter, the actual procedure that would be called when running Proc would be determined at run time, and it might be different for every call to Proc. In this case, we wouldn't be able to determine at compilation time which procedure would be called in every case. (In some cases, however, it could still be possible to determine which procedure is called by analyzing the argument that is passed to Proc.)

15.15.2 Access to subprogram declaration

We declare an access to a subprogram as a type by writing **access procedure** or **access function** and the corresponding prototype: Listing 129: access_to_subprogram_types.ads

```
package Access_To_Subprogram_Types is

type Access_To_Procedure is
access procedure (I : in out Integer);

type Access_To_Function is
access function (I : Integer) return Integer;
end Access To Subprogram Types;
```

Code block metadata

In the designated profile of the access type declarations, we list all the parameters that we expect in the subprogram.

We can use those types to declare access to subprograms — as subprogram parameters, for example:

Listing 130: access_to_subprogram_params.ads

```
with Access_To_Subprogram_Types;
use Access_To_Subprogram_Types;
package Access_To_Subprogram_Params is
procedure Proc (P : Access_To_Procedure);
end Access_To_Subprogram_Params;
```

Listing 131: access_to_subprogram_params.adb

```
package body Access_To_Subprogram_Params is
1
2
      procedure Proc (P : Access_To_Procedure) is
3
         I : Integer := 0;
4
      begin
5
         P (I);
6
7
         -- P.all (I);
      end Proc;
8
  end Access_To_Subprogram_Params;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Access_To_Subprogram_Types

MD5: 17c1a07f48d9fb0efef37aa4c5ec8a51
```

In the implementation of the Proc procedure of the code example, we call the P procedure by simply passing I as a parameter. In this case, P is automatically dereferenced. We may, however, explicitly dereference P by writing P.all (I).

Before we use this package, let's implement a simple procedure that we'll use later on:

Listing 132: add_ten.ads

```
procedure Add_Ten (I : in out Integer);
```

Listing 133: add_ten.adb

```
1 procedure Add_Ten (I : in out Integer) is
2 begin
3 I := I + 10;
4 end Add_Ten;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Access_To_Subprogram_Types

MD5: 8553ad7329bf1ed727147b47b7355a70
```

Now, we can get access to a subprogram by using the **Access** attribute and pass it as an actual parameter:

Listing 134: show access to subprograms.adb

```
with Access_To_Subprogram_Params;
1
   use Access_To_Subprogram_Params;
2
3
   with Add_Ten;
4
5
   procedure Show_Access_To_Subprograms is
6
   begin
7
      Proc (Add_Ten'Access);
8
                      `Getting access to Add Ten
9
                       procedure and passing it
10
                       to Proc
11
   end Show_Access_To_Subprograms;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Access_To_Subprogram_Types

MD5: 599e9d1306da48e3c532692b34c02a1d
```

Here, we get access to the Add_Ten procedure and pass it to the Proc procedure.

In the Ada Reference Manual

3.10 Access Types²⁸⁹

15.15.3 Objects of access-to-subprogram type

In the previous example, the Proc procedure had a parameter of access-to-subprogram type. In addition to parameters, we can of course declare *objects* of access-to-subprogram types as well. For example, we can extend our previous test application and declare an object P of access-to-subprogram type. Before we do so, however, let's implement another small procedure that we'll use later on:

²⁸⁹ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

Listing 135: add_twenty.ads

```
procedure Add_Twenty (I : in out Integer);
```

Listing 136: add_twenty.adb

```
procedure Add_Twenty (I : in out Integer) is
begin
I := I + 20;
end Add_Twenty;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Access_To_Subprogram_Types

MD5: 697959b806f6f2bfba248ec15c47883b
```

In addition to Add_Ten, we've implemented the Add_Twenty procedure, which we use in our extended test application:

Listing 137: show access to subprograms.adb

```
with Access_To_Subprogram_Types;
1
   use Access_To_Subprogram_Types;
2
3
   with Access_To_Subprogram_Params;
4
   use Access_To_Subprogram_Params;
5
6
   with Add_Ten;
7
   with Add_Twenty;
8
9
   procedure Show_Access_To_Subprograms is
10
      Ρ
                : Access_To_Procedure;
11
      Some_Int : Integer := 0;
12
   begin
13
      P := Add_Ten'Access;
14
                     ^ Getting access to Add_Ten
15
       - -
       - -
                       procedure and assigning it
16
       - -
                       to P
17
18
      Proc (P);
19
             ^ Passing access-to-subprogram as an
       - -
20
              actual parameter
       - -
21
22
      P (Some Int);
23
       -- ^ Using access-to-subprogram object in a
24
       - -
           subprogram call
25
26
      P := Add_Twenty'Access;
27
                        ^ Getting access to Add Twenty
28
       - -
                          procedure and assigning it
       - -
29
                          to P
       - -
30
31
      Proc (P);
32
      P (Some Int);
33
   end Show Access To Subprograms;
34
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Access_To_Subprogram_Types
```

```
MD5: 7b4ea19187806e88ba65847876cafb4f
```

In the Show_Access_To_Subprograms procedure, we see the declaration of our access-tosubprogram object P (of Access_To_Procedure type). We get access to the Add_Ten procedure and assign it to P, and we then do the same for the Add Twenty procedure.

We can use an access-to-subprogram object either as the actual parameter of a subprogram call, or in a subprogram call. In the code example, we're passing P as the actual parameter of the Proc procedure in the Proc (P) calls. Also, we're calling the subprogram assigned to (designated by the current value of) P in the P (Some_Int) calls.

15.15.4 Components of access-to-subprogram type

In addition to declaring subprogram parameters and objects of access-to-subprogram types, we can declare components of these types. For example:

```
Listing 138: access_to_subprogram_types.ads
```

```
package Access_To_Subprogram_Types is
1
2
      type Access_To_Procedure is
3
        access procedure (I : in out Integer);
4
5
      type Access_To_Function is
6
         access function (I : Integer) return Integer;
7
8
      type Access_To_Procedure_Array is
9
         array (Positive range <>) of
10
           Access_To_Procedure;
11
12
      type Access To Function Array is
13
        array (Positive range <>) of
14
           Access_To_Function;
15
16
      type Rec Access To Procedure is record
17
         AP : Access_To_Procedure;
18
      end record;
19
20
      type Rec_Access_To_Function is record
21
         AF : Access_To_Function;
22
      end record;
23
24
   end Access_To_Subprogram_Types;
25
```

Code block metadata

Here, the access-to-procedure type Access_To_Procedure is used as a component of the array type Access_To_Procedure_Array and the record type Rec_Access_To_Procedure. Similarly, the access-to-function type Access_To_Function type is used as a component of the array type Access_To_Function_Array and the record type Rec_Access_To_Function.

Let's see two test applications using these types. First, let's use the Access_To_Procedure_Array array type in a test application:

```
Listing 139: show_access_to_subprograms.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_To_Subprogram_Types;
3
   use Access_To_Subprogram_Types;
4
5
   with Add Ten;
6
   with Add Twenty;
7
8
   procedure Show_Access_To_Subprograms is
9
      PA : constant
10
              Access_To_Procedure_Array (1 .. 2) :=
11
                (Add_Ten'Access,
12
                 Add_Twenty'Access);
13
14
      Some Int : Integer := 0;
15
   begin
16
      Put Line ("Some Int: " & Some Int'Image);
17
18
      for I in PA'Range loop
19
         PA (I) (Some_Int);
20
          Put Line ("Some Int: " & Some Int'Image);
21
      end loop;
22
   end Show_Access_To_Subprograms;
23
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_ ⇔Subprograms.Access_To_Subprogram_Types MD5: f1d10056b4b3424bd30d954f34caa255

Runtime output

Some_Int: 0 Some_Int: 10 Some_Int: 30

Here, we declare the PA array and use the access to the Add_Ten and Add_Twenty procedures as its components. We can call any of these procedures by simply specifying the index of the component, e.g. PA (2). Once we specify the procedure we want to use, we simply pass the parameters, e.g.: PA (2) (Some_Int).

Now, let's use the Rec_Access_To_Procedure record type in a test application:

Listing 140: show_access_to_subprograms.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_To_Subprogram_Types;
3
   use Access_To_Subprogram_Types;
4
5
   with Add Ten;
6
   with Add_Twenty;
7
8
   procedure Show_Access_To_Subprograms is
9
               : Rec Access To Procedure;
      RA
10
      Some_Int : Integer := 0;
11
   begin
12
      Put_Line ("Some_Int: " & Some_Int'Image);
13
14
      RA := (AP => Add_Ten'Access);
15
```

```
16 RA.AP (Some_Int);
17 Put_Line ("Some_Int: " & Some_Int'Image);
18
19 RA := (AP => Add_Twenty'Access);
20 RA.AP (Some_Int);
21 Put_Line ("Some_Int: " & Some_Int'Image);
22 end Show_Access_To_Subprograms;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Access_To_Subprogram_Types

MD5: 4b23b5f6a8c252a1a014a2b54fa32c1a
```

Runtime output

Some_Int: 0 Some_Int: 10 Some_Int: 30

Here, we declare two record aggregates where we specify the AP component, e.g.: (AP => Add_Ten'Access), which indicates the access-to-subprogram we want to use. We can call the subprogram by simply accessing the AP component, i.e.: RA.AP.

15.15.5 Access-to-subprogram as discriminant types

As you might expect, we can use access-to-subprogram types when declaring discriminants. In fact, when we were talking about *discriminants as access values* (page 603) earlier on, we used access-to-object types in our code examples, but we could have used access-to-subprogram types as well. For example:

Listing 141: custom_processing.ads

```
package Custom Processing is
1
2
       -- Declaring an access type:
3
      type Integer_Processing is
4
        access procedure (I : in out Integer);
5
6
       -- Declaring a discriminant with this
7
       -- access type:
8
       type Rec (IP : Integer_Processing) is
9
         private:
10
11
      procedure Init (R
                          : in out Rec;
12
                        Value :
                                        Integer);
13
14
      procedure Process (R : in out Rec);
15
16
      procedure Show (R : Rec);
17
18
   private
19
20
       type Rec (IP : Integer Processing) is
21
       record
22
          I : Integer := 0;
23
       end record;
24
25
   end Custom Processing;
26
```

```
Listing 142: custom_processing.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Custom Processing is
3
4
      procedure Init (R
                             : in out Rec;
5
                        Value :
                                   Integer) is
6
      begin
7
         R.I := Value;
8
      end Init;
9
10
      procedure Process (R : in out Rec) is
11
      begin
12
          R.IP (R.I);
13
              ~~~~
          - -
14
             Calling procedure that we specified as
          - -
15
          -- the record's discriminant
16
       end Process;
17
18
      procedure Show (R : Rec) is
19
      begin
20
          Put Line ("R.I = "
21
                    & Integer'Image (R.I));
22
      end Show;
23
24
   end Custom_Processing;
25
```

Code block metadata

In this example, we declare the access-to-subprogram type Integer_Processing, which we use as the IP discriminant of the Rec type. In the Process procedure, we call the IP procedure that we specified as the record's discriminant (R.IP (R.I)).

Before we look at a test application for this package, let's implement another small procedure:

Listing 143: mult_two.ads

```
procedure Mult_Two (I : in out Integer);
```

Listing 144: mult_two.adb

```
1 procedure Mult_Two (I : in out Integer) is
2 begin
3 I := I * 2;
4 end Mult_Two;
```

Code block metadata

Now, let's look at the test application:

```
with Ada.Text IO;
                             use Ada.Text I0;
1
2
   with Custom_Processing; use Custom_Processing;
3
4
   with Add Ten:
5
   with Mult Two;
6
7
   procedure Show Access To Subprogram Discriminants
8
   is
9
10
      R Add Ten : Rec (IP => Add Ten'Access);
11
12
                Using access-to-subprogram as a
       - -
13
                discriminant
       - -
14
15
      R Mult Two : Rec (IP => Mult Two'Access);
16
                           ^^^^^
17
                Using access-to-subprogram as a
18
       - -
                discriminant
19
20
   begin
21
      Init (R_Add_Ten, 1);
22
      Init (R_Mult_Two, 2);
23
24
      Put Line ("---- R Add Ten ----");
25
      Show (R_Add_Ten);
26
27
      Put_Line ("Calling Process procedure...");
28
      Process (R Add Ten);
29
      Show (R_Add_Ten);
30
31
      Put Line ("---- R Mult Two ----");
32
      Show (R_Mult_Two);
33
34
      Put Line ("Calling Process procedure...");
35
      Process (R Mult Two);
36
      Show (R Mult Two);
37
   end Show Access To Subprogram Discriminants;
38
```

Listing 145: show_access_to_subprogram_discriminants.adb

Code block metadata

Runtime output

```
---- R_Add_Ten ----
R.I = 1
Calling Process procedure...
R.I = 11
---- R_Mult_Two ----
R.I = 2
Calling Process procedure...
R.I = 4
```

In this procedure, we declare the R_Add_Ten and R_Mult_Two of Rec type and specify the access to Add_Ten and Mult_Two, respectively, as the IP discriminant. The procedure we specified here is then called inside a call to the Process procedure.

15.15.6 Access-to-subprograms as formal parameters

We can use access-to-subprograms types when declaring formal parameters. For example, let's revisit the Custom_Processing package from the previous section and convert it into a generic package.

```
Listing 146: gen_custom_processing.ads
```

```
generic
1
      type T is private;
2
3
4
       -- Declaring formal access-to-subprogram
5
       - -
          type:
6
7
       type T Processing is
8
        access procedure (Element : in out T);
9
10
11
       -- Declaring formal access-to-subprogram
12
       -- parameter:
13
14
      Proc : T_Processing;
15
16
      with function Image_T (Element : T)
17
                               return String;
18
   package Gen Custom Processing is
19
20
      type Rec is private;
21
22
      procedure Init (R : in out Rec;
23
                        Value :
                                       T):
24
25
      procedure Process (R : in out Rec);
26
27
      procedure Show (R : Rec);
28
29
   private
30
31
       type Rec is record
32
         Comp : T;
33
      end record;
34
35
   end Gen_Custom_Processing;
36
```

Listing 147: gen_custom_processing.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Gen_Custom_Processing is
3
4
      procedure Init (R : in out Rec;
5
                       Value :
                                       T) is
6
      begin
7
         R.Comp := Value;
8
      end Init;
9
10
      procedure Process (R : in out Rec) is
11
      begin
12
         Proc (R.Comp);
13
      end Process;
14
15
```

```
16 procedure Show (R : Rec) is
17 begin
18 Put_Line ("R.Comp = "
19 & Image_T (R.Comp));
20 end Show;
21
22 end Gen_Custom_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Access_To_Subprogram_Types

MD5: 6f06e066bafa5f02abb3ee1b33ea0831
```

In this version of the procedure, instead of declaring Proc as a discriminant of the Rec record, we're declaring it as a formal parameter of the Gen_Custom_Processing package. Also, we're declaring an access-to-subprogram type (T_Processing) as a formal parameter. (Note that, in contrast to these two parameters that we've just mentioned, Image_T is not a formal access-to-subprogram parameter: it's actually just a formal subprogram.)

We then instantiate the Gen_Custom_Processing package in our test application:

Listing 148: show_access_to_subprogram_as_formal_parameter.adb

```
with Gen Custom Processing;
1
2
   with Add Ten;
3
4
   with Ada.Text IO; use Ada.Text IO;
5
6
   procedure
7
     Show Access To Subprogram As Formal Parameter
8
9
   is
       type Integer Processing is
10
         access procedure (I : in out Integer);
11
12
      package Custom Processing is new
13
         Gen Custom Processing
14
           (T
                         => Integer,
15
            T_Processing => Integer_Processing,
16
            - -
17
            - -
                             access-to-subprogram type
18
                          => Add Ten'Access,
            Proc
19
                              ~~~~~~~~
            - -
20
            - -
                             access-to-subprogram
21
            Image T
                         => Integer'Image);
22
      use Custom Processing;
23
24
      R_Add_Ten : Rec;
25
26
   begin
27
      Init (R_Add_Ten, 1);
28
29
      Put Line ("---- R Add Ten ----");
30
      Show (R Add Ten);
31
32
      Put Line ("Calling Process procedure...");
33
      Process (R Add Ten);
34
       Show (R Add Ten);
35
   end Show_Access_To_Subprogram_As_Formal_Parameter;
36
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_ Subprograms.Access_To_Subprogram_Types MD5: 6ae27ebd59e5307551e9a38f3b94c70c

Runtime output

---- R_Add_Ten ----R.Comp = 1 Calling Process procedure... R.Comp = 11

Here, we instantiate the Gen_Custom_Processing package as Custom_Processing and specify the access-to-subprogram type and the access-to-subprogram.

15.15.7 Selecting subprograms

A practical application of access to subprograms is that it enables us to dynamically select a subprogram and pass it to another subprogram, where it can then be called.

For example, we may have a Process procedure that receives a logging procedure as a parameter (Log_Proc). Also, this parameter may be **null** by default — so that no procedure is called if the parameter isn't specified:

Listing 149: data processing.ads

```
package Data Processing is
1
2
      type Data_Container is
3
        array (Positive range <>) of Float;
4
5
      type Log Procedure is
6
        access procedure (D : Data_Container);
7
8
      procedure Process
9
              : in out Data_Container;
        (D
10
         Log Proc :
                       Log Procedure := null);
11
12
   end Data Processing;
13
```

Listing 150: data_processing.adb

```
package body Data_Processing is
1
2
      procedure Process
3
        (D : in out Data_Container;
4
         Log Proc : Log Procedure := null) is
5
      begin
6
         -- missing processing part...
7
8
         if Log_Proc /= null then
9
            Log Proc (D);
10
         end if;
11
      end Process;
12
13
   end Data_Processing;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Log_Procedure

MD5: 59399e0809deb476f608faab7e4398bd
```

In the implementation of Process, we check whether Log_Proc is null or not. (If it's not null, we call the procedure. Otherwise, we just skip the call.)

Now, let's implement two logging procedures that match the expected form of the Log_Procedure type:

```
Listing 151: log_element_per_line.adb
```

```
with Ada.Text IO;
                           use Ada.Text IO;
1
   with Data_Processing; use Data_Processing;
2
   procedure Log_Element_Per_Line
4
     (D : Data Container) is
5
   begin
6
      Put Line ("Elements: ");
7
      for V of D loop
8
         Put Line (V'Image);
9
      end loop;
10
      Put_Line ("-----");
11
   end Log_Element_Per_Line;
12
```

Listing 152: log csv.adb

```
with Ada.Text_I0;
                          use Ada.Text_I0;
1
   with Data_Processing; use Data_Processing;
2
3
   procedure Log_Csv (D : Data_Container) is
4
5
   begin
      for I in D'First .. D'Last - 1 loop
6
         Put (D (I)'Image & ", ");
7
      end loop;
8
      Put (D (D'Last)'Image);
9
      New Line;
10
  end Log Csv;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Log_Procedure

MD5: 468789f7331ffcd16f754f7116b076d7
```

Finally, we implement a test application that selects each of the logging procedures that we've just implemented:

```
Listing 153: show_access_to_subprograms.adb
```

```
with Ada.Text_I0;
                          use Ada.Text_I0;
1
   with Data_Processing; use Data_Processing;
2
3
   with Log_Element_Per_Line;
4
   with Log_Csv;
5
6
   procedure Show_Access_To_Subprograms is
7
      D : Data_Container (1 .. 5) := (others => 1.0);
8
   begin
9
      Put_Line ("==== Log_Element_Per_Line ====");
10
      Process (D, Log_Element_Per_Line'Access);
11
12
      Put Line ("==== Log Csv ====");
13
      Process (D, Log_Csv'Access);
14
15
      Put Line ("==== None ====");
16
```

17 Process (D); 18 end Show Access To Subprograms;

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Log_Procedure

MD5: 134aa682cea1999efa0ea97052f315c8
```

Runtime output

```
==== Log_Element_Per_Line ====
Elements:
   1.00000E+00
   1.00000E+00
   1.00000E+00
   1.00000E+00
   1.00000E+00
   1.00000E+00
   1.00000E+00, 1.00000E+00, 1.00000E+00, 1.00000E+00
==== None ====
```

Here, we use the **Access** attribute to get access to the Log_Element_Per_Line and Log_Csv procedures. Also, in the third call, we don't pass any access as an argument, which is then **null** by default.

15.15.8 Null exclusion

We can use null exclusion when declaring an access to subprograms. By doing so, we ensure that a subprogram must be specified — either as a parameter or when initializing an access object. Otherwise, an exception is raised. Let's adapt the previous example and introduce the Init_Function type:

Listing 154: data_processing.ads

```
package Data_Processing is
1
2
      type Data Container is
3
        array (Positive range <>) of Float;
4
5
      type Init Function is
6
         not null access function return Float;
7
8
      procedure Process
9
                  : in out Data Container;
         (D
10
         Init Func :
                              Init Function);
11
12
   end Data_Processing;
13
```

Listing 155: data_processing.adb

```
1 package body Data_Processing is
2
3 procedure Process
4 (D : in out Data_Container;
5 Init_Func : Init_Function) is
6 begin
7 for I in D'Range loop
```

```
B D (I) := Init_Func.all;
P end loop;
B end Process;
B end Data_Processing;
```

In this case, we specify that Init_Function is **not null access** because we want to always be able to call this function in the Process procedure (i.e. without raising an exception).

When an access to a subprogram doesn't have parameters — which is the case for the subprograms of Init_Function type — we need to explicitly dereference it by writing . **all**. (In this case, .**all** isn't optional.) Therefore, we have to write Init_Func.**all** in the implementation of the Process procedure of the code example.

Now, let's declare two simple functions — Init_Zero and Init_One — that return 0.0 and 1.0, respectively:

Listing 156: init_zero.ads

1 function Init_Zero return Float;

Listing 157: init_one.ads

1 function Init_One return Float;

Listing 158: init_zero.adb

- 1 function Init_Zero return Float is
- 2 begin
- 3 return 0.0;
- 4 end Init_Zero;

Listing 159: init one.adb

- 1 function Init_One return Float is
- 2 begin
- 3 return 1.0;
- 4 end Init_One;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_ →Subprograms.Access_Init_Function MD5: 444110d50ddb430fd5be31cf1b417fc8

Finally, let's see a test application where we select each of the init functions we've just implemented:

Listing 160: log element per line.adb

```
with Ada.Text IO;
                          use Ada.Text I0;
1
   with Data Processing; use Data Processing;
2
3
   procedure Log Element Per Line
4
     (D : Data Container) is
5
   begin
6
      Put Line ("Elements: ");
7
      for V of D loop
8
         Put Line (V'Image);
9
      end loop;
10
```

Put_Line ("-----");
Put_Line [Put_Line;

Listing 161: show access to subprograms.adb

```
with Ada.Text IO;
                           use Ada.Text I0;
1
   with Data Processing; use Data Processing;
2
3
   with Init Zero;
4
   with Init_One;
5
6
   with Log_Element_Per_Line;
7
8
   procedure Show Access To Subprograms is
9
      D : Data Container (1 .. 5) := (others => 1.0);
10
   begin
11
      Put Line ("==== Init Zero ====");
12
      Process (D, Init Zero'Access);
13
      Log_Element_Per_Line (D);
14
15
      Put_Line ("==== Init_One ====");
16
      Process (D, Init_One'Access);
17
      Log_Element_Per_Line (D);
18
19
          Put Line ("==== None ====");
       - -
20
       -- Process (D, null);
21
          Log Element Per Line (D);
22
   end Show Access To Subprograms;
23
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Access_Init_Function

MD5: ae0e3fd58e9bb83061248967c709190a
```

Runtime output

```
=== Init_Zero ====
Elements:
    0.00000E+00
    0.00000E+00
    0.00000E+00
    0.00000E+00
-----
==== Init_One ====
Elements:
    1.00000E+00
    1.00000E+00
    1.00000E+00
    1.00000E+00
    1.00000E+00
```

Here, we use the **Access** attribute to get access to the Init_Zero and Init_One functions. Also, if we uncomment the call to Process with **null** as an argument for the init function, we see that the Constraint_Error exception is raised at run time — as the argument cannot be **null** due to the null exclusion.

E	Note
Т	his example was originally written by Robert A. Duff and was part of the Gem $#24^{29}$
Чe	re's another example, first with null :
	Listing 162: show_null_procedure.ads
pad	<pre>ckage Show_Null_Procedure is type Element is limited null record; Not implemented yet</pre>
	<pre>type Ref_Element is access all Element;</pre>
	<pre>type Table is limited null record; Not implemented yet</pre>
	<pre>type Iterate_Action is access procedure (X : not null Ref_Element);</pre>
	<pre>procedure Iterate (T : Table; Action : Iterate_Action := null); If Action is null, do nothing.</pre>
enc	d Show_Null_Procedure;
Co	de block metadata
ے ج	oject: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_ Subprograms.Null_Procedure 5: ac21dd76ed9fb7f26839c24210cf4425

and without **null**:

Listing 163: show null procedure.ads

```
package Show_Null_Procedure is
  type Element is limited null record;
  -- Not implemented yet
  type Ref_Element is access all Element;
  type Table is limited null record;
  -- Not implemented yet
  procedure Do_Nothing
  (X : not null Ref_Element) is null;
  type Iterate_Action is
   access procedure
    (X : not null Ref_Element);
  procedure Iterate
   (T : Table;
    Action : not null Iterate_Action
```

22

1

2

3 4

5 6

7

8 9

10

end Show Null Procedure;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_ →Subprograms.Null_Procedure MD5: 7341d8f23cd4efe45698481be452a9e8

:= Do Nothing'Access);

The style of the second Iterate is clearly better because it makes use of the syntax to indicate that a procedure is expected. This is a complete package that includes both versions of the Iterate procedure:

Listing 164: example.ads

```
package Example is
1
2
       type Element is limited private;
3
       type Ref Element is access all Element;
4
5
       type Table is limited private;
6
7
      type Iterate Action is
8
         access procedure
9
           (X : not null Ref Element);
10
11
      procedure Iterate
12
         (T : Table;
13
          Action : Iterate Action := null);
14
       -- If Action is null, do nothing.
15
16
       procedure Do Nothing
17
        (X : not null Ref_Element) is null;
18
      procedure Iterate 2
19
         (T : Table;
20
          Action : not null Iterate Action
21
                      := Do Nothing'Access);
22
23
   private
24
      type Element is limited
25
          record
26
             Component : Integer;
27
          end record;
28
      type Table is limited null record;
29
30 6944 Example;
```

```
Listing 165: example.adb
   package body Example is
1
2
       An_Element : aliased Element;
3
4
       procedure Iterate
5
         (T : Table;
6
          Action : Iterate Action := null)
7
       is
8
       begin
9
          if Action /= null then
10
             Action (An Element'Access);
11
             -- In a real program, this would do
12
                 something more sensible.
13
          end if;
14
       end Iterate;
15
16
       procedure Iterate 2
17
         (T : Table;
18
19
          Action : not null Iterate Action
                      := Do Nothing'Access)
20
       is
21
       begin
22
          Action (An_Element'Access);
23
          -- In a real program, this would do
24
          -- something more sensible.
25
       end Iterate 2;
26
27
   end Example;
28
```

Listing 166: show_example.adb

```
with Example; use Example;
procedure Show_Example is
T : Table;
begin
Iterate_2 (T);
r end Show_Example;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_ →Subprograms.Complete_Not_Null_Procedure MD5: ab0a41e0d39a8a16b0b69f8c6b2a43fd

Writing **not null** Iterate_Action might look a bit more complicated, but it's worthwhile, and anyway, as mentioned earlier, the compatibility requirement requires that the **not null** be explicit, rather than the other way around.

15.15.9 Access to protected subprograms

Up to this point, we've discussed access to *normal* Ada subprograms. In some situations, however, we might want to have access to protected subprograms. To do this, we can simply declare a type using **access protected**:

²⁹⁰ https://www.adacore.com/gems/ada-gem-24

Listing 167:	simple	protected	access.ads

```
package Simple Protected Access is
1
2
      type Access_Proc is
3
         access protected procedure;
4
5
      protected Obj is
6
7
          procedure Do Something;
8
9
      end Obj;
10
11
      Acc : Access_Proc := Obj.Do_Something'Access;
12
13
   end Simple_Protected_Access;
14
```

Listing 168: simple_protected_access.adb

```
package body Simple_Protected_Access is
1
2
      protected body Obj is
3
4
          procedure Do_Something is
5
          begin
6
             - -
                 Not doing anything
7
             - -
                 for the moment...
8
             null;
9
          end Do_Something;
10
11
      end Obj;
12
13
   end Simple_Protected_Access;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

⇔Subprograms.Simple_Protected_Access

MD5: d82f7c90355e9810bd1e35f65e278626
```

Here, we declare the Access_Proc type as an access type to protected procedures. Then, we declare the variable Acc and assign to it the access to the Do_Something procedure (of the protected object Obj).

Now, let's discuss a more useful example: a simple system that allows us to register protected procedures and execute them. This is implemented in Work_Registry package:

```
Listing 169: work_registry.ads
```

```
package Work_Registry is
1
2
      type Work_Id is tagged limited private;
3
4
      type Work Handler is
5
        access protected procedure (T : Work_Id);
6
7
      subtype Valid Work Handler is
8
        not null Work_Handler;
9
10
      type Work_Handlers is
11
         array (Positive range <>) of Work_Handler;
12
```

```
13
       protected type Work_Handler_Registry
14
         (Last : Positive)
15
       is
16
17
          procedure Register (T : Valid_Work_Handler);
18
19
          procedure Reset;
20
21
          procedure Process_All;
22
23
       private
24
25
          D
                : Work_Handlers (1 .. Last);
26
          Curr : Natural := 0;
27
28
       end Work_Handler_Registry;
29
30
   private
31
32
       type Work_Id is tagged limited null record;
33
34
   end Work_Registry;
35
```

Listing 170: work_registry.adb

```
package body Work_Registry is
1
2
       protected body Work_Handler_Registry is
3
4
          procedure Register (T : Valid_Work_Handler)
5
          is
6
          begin
7
             if Curr < Last then</pre>
8
                 Curr := Curr + 1;
9
                 D (Curr) := T;
10
             end if;
11
          end Register;
12
13
          procedure Reset is
14
          begin
15
             Curr := 0;
16
          end Reset;
17
18
          procedure Process All is
19
             Dummy_ID : Work_Id;
20
          begin
21
              for I in D'First .. Curr loop
22
                 D (I).all (Dummy ID);
23
              end loop;
24
          end Process_All;
25
26
       end Work_Handler_Registry;
27
28
   end Work_Registry;
29
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Protected_Access_Init_Function

MD5: 5dfa8ab098900ab4f6b7575e1cde5e53
```

Here, we declare the protected Work_Handler_Registry type with the following subprograms:

- Register, which we can use to register a protected procedure;
- Reset, which we can use to reset the system; and
- Process_All, which we can use to call all procedures that were registered in the system.

Work_Handler is our access to protected subprogram type. Also, we declare the Valid_Work_Handler subtype, which excludes **null**. By doing so, we can ensure that only valid procedures are passed to the Register procedure. In the protected Work_Handler_Registry type, we store the procedures in an array (of Work_Handlers type).

Important

Note that, in the type declaration Work_Handler, we say that the protected procedure must have a parameter of Work_Id type. In this example, this parameter is just used to *bind* the procedure to the Work_Handler_Registry type. The Work_Id type itself is actually declared as a null record (in the private part of the package), and it isn't really useful on its own.

If we had declared **type Work_Handler is access protected procedure**; instead, we would be able to register *any* protected procedure into the system, even the ones that might not be suitable for the system. By using a parameter of Work_Id type, however, we make use of strong typing to ensure that only procedures that were designed for the system can be registered.

In the next part of the code, we declare the Integer_Storage type, which is a simple protected type that we use to store an integer value:

Listing 171: integer_storage_system.ads

```
with Work_Registry;
1
2
   package Integer_Storage_System is
3
4
      protected type Integer_Storage is
5
6
          procedure Set (V : Integer);
7
8
          procedure Show (T : Work_Registry.Work_Id);
9
10
      private
11
12
          I : Integer := 0;
13
14
      end Integer_Storage;
15
16
      type Integer_Storage_Access is
17
         access Integer_Storage;
18
19
      type Integer Storage Array is
20
         array (Positive range <>) of
21
           Integer_Storage_Access;
22
23
   end Integer_Storage_System;
24
```

Listing 172: integer_storage_system.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Integer_Storage_System is
3
4
      protected body Integer Storage is
5
6
          procedure Set (V : Integer) is
7
          begin
8
             I := V;
9
          end Set;
10
11
          procedure Show (T : Work_Registry.Work_Id)
12
          is
13
             pragma Unreferenced (T);
14
          begin
15
             Put Line ("Value: " & Integer'Image (I));
16
          end Show;
17
18
       end Integer_Storage;
19
20
   end Integer_Storage_System;
21
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Protected_Access_Init_Function

MD5: a388d792bc85709785d324c914d9d236
```

For the Integer_Storage type, we declare two procedures:

- · Set, which we use to assign a value to the (protected) integer value; and
- Show, which we use to show the integer value that is stored in the protected object.

The Show procedure has a parameter of Work_Id type, which indicates that this procedure was designed to be registered in the system of Work_Handler_Registry type.

Finally, we have a test application in which we declare a registry (WHR) and an array of "protected integer objects" (Int_Stor):

Listing 173: show_access_to_protected_subprograms.adb

```
with Work_Registry;
1
   use Work_Registry;
2
3
   with Integer_Storage_System;
4
   use Integer_Storage_System;
5
6
   procedure Show Access To Protected Subprograms is
7
8
                : Work_Handler_Registry (5);
      WHR
9
      Int_Stor : Integer_Storage_Array (1 .. 3);
10
11
   begin
12
      -- Allocate and initialize integer storage
13
14
      -- (For the initialization, we're just
15
       -- assigning the index here, but we could
16
      -- really have used any integer value.)
17
18
      for I in Int_Stor'Range loop
19
```

```
Int_Stor (I) := new Integer_Storage;
20
          Int_Stor (I).Set (I);
21
      end loop;
22
23
       -- Register handlers
24
25
      for I in Int_Stor'Range loop
26
         WHR.Register (Int_Stor (I).all.Show'Access);
27
      end loop;
28
29
       -- Now, use Process All to call the handlers
30
          (in this case, the Show procedure for
31
           each protected object from Int_Stor).
       - - -
32
33
      WHR.Process_All;
34
35
   end Show_Access_To_Protected_Subprograms;
36
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_To_

→Subprograms.Protected_Access_Init_Function

MD5: 44c24ef07333e1d31844cc2ea6d91ab6
```

Runtime output

Value: 1 Value: 2 Value: 3

The work handler registry (WHR) has a maximum capacity of five procedures, whereas the Int_Stor array has a capacity of three elements. By calling WHR.Register and passing Int_Stor (I).all.Show'Access, we register the Show procedure of each protected object from Int_Stor.

1 Important

Note that the components of the Int_Stor array are of Integer_Storage_Access type, which is declared as an access to Integer_Storage objects. Therefore, we have to dereference the object (by writing Int_Stor (I).all) before getting access to the Show procedure (by writing .Show'Access).

We have to use an access type here because we cannot pass the access (to the Show procedure) of a local object in the call to the Register procedure. Therefore, the protected objects (of Integer_Storage type) cannot be local.

This issue becomes evident if we replace the declaration of Int_Stor with a local array (and then adapt the remaining code). If we do this, we get a compilation error in the call to Register:

Listing 174: show_access_to_protected_subprograms.adb

```
with Work_Registry;
1
  use Work_Registry;
2
3
  with Integer_Storage_System;
4
  use Integer_Storage_System;
5
6
  procedure Show Access To Protected Subprograms
7
8
  is
      WHR
                : Work_Handler_Registry (5);
9
```

```
10
       Int_Stor : array (1 .. 3) of Integer_Storage;
11
12
   begin
13
          Allocate and initialize integer storage
       - -
14
15
       - -
           (For the initialization, we're just
16
       - -
           assigning the index here, but we could
17
       - -
          really have used any integer value.)
18
19
       for I in Int Stor'Range loop
20
          -- Int Stor (I) := new Integer Storage;
21
          Int_Stor (I).Set (I);
22
       end loop;
23
24
       -- Register handlers
25
26
       for I in Int Stor'Range loop
27
          WHR.Register (Int_Stor (I).Show'Access);
28
                          ^ ERROR!
29
       end loop;
30
31
       -- Now, call the handlers
32
       -- (i.e. the Show procedure of each
33
            protected object).
       - -
34
35
       WHR.Process All;
36
37
   end Show_Access_To_Protected_Subprograms;
38
```

Code block metadata

Build output

As we've just discussed, this error is due to the fact that Int_Stor is now a "local" protected object, and the accessibility rules don't allow mixing it with non-local accesses in order to prevent the possibility of dangling references.

When we call WHR.Process_All, the registry system calls each procedure that has been registered with the system. When looking at the values displayed by the test application, we may notice that each call to Show is referring to a different protected object. In fact, even though we're passing just the access to a protected *procedure* in the call to Register, that access is also associated to a specific protected object. (This is different from access to non-protected subprograms we've discussed previously: in that case, there's no object associated.) If we replace the argument to Register by Int_Stor (2).all.Show'Access, for example, the three Show procedures registered in the system will now refer to the same protected object (stored at Int Stor (2)).

Also, even though we have registered the same procedure (Show) of the same type (Integer_Storage) in all calls to Register, we could have used a different protected procedure — and of a different protected type. As an exercise, we could, for example, create a new type called Float_Storage (based on the code that we used for the Integer_Storage type) and register some objects of Float_Storage type into the system (with a couple of additional calls to Register). If we then call WHR.Process_All, we'd see that the system is

able to cope with objects of both Integer_Storage and Float_Storage types. In fact, the system implemented with the Work_Handler_Registry can be seen as "type agnostic," as it doesn't care about which type the protected objects have — as long as the subprograms we want to register are conformant to the Valid_Work_Handler type.

15.16 Accessibility Rules and Access-To-Subprograms

In general, the accessibility rules that we discussed *previously for access-to-objects* (page 645) also apply to access-to-subprograms. In this section, we discuss minor differences when applying those rules to access-to-subprograms.

In our discussion about accessibility rules, we've looked into *accessibility levels* (page 646) and the *accessibility rules* (page 647) that are based on those levels. The same accessibility rules apply to access-to-subprograms. *As we said previously* (page 649), operations targeting objects at a *less-deep* level are illegal, as it's the case for subprograms as well:

Listing 175: access_to_subprogram_types.ads

```
package Access_To_Subprogram_Types is

type Access_To_Procedure is
access procedure (I : in out Integer);

type Access_To_Function is
access function (I : Integer) return Integer;
end Access To Subprogram Types;
```

Listing 176: show_access_to_subprogram_error.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Access_To_Subprogram_Types;
3
   use Access To Subprogram Types;
4
5
   procedure Show Access To Subprogram Error is
6
      Func : Access To Function;
7
8
      Value : Integer := 0;
9
   beain
10
      declare
11
          function Add One (I : Integer)
12
                             return Integer is
13
            (I + 1);
14
      begin
15
         Func := Add One'Access;
16
          -- This assignment is illegal because the
17
          -- Access To Function type is less deep
18
         -- than Add One.
19
      end;
20
21
      Put Line ("Value: " & Value'Image);
22
      Value := Func (Value);
23
      Put Line ("Value: " & Value'Image);
24
   end Show Access To Subprogram Error;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Accessibility_Rules_

⇔Access_To_Subprograms.Access_To_Subprogram_Accessibility_Error_Less_Deep

(continues on next page)
```

```
MD5: 2a068732606a1fee156e82515febe9c4
```

Build output

```
show_access_to_subprogram_error.adb:16:15: error: subprogram must not be deeper_

→than access type

gprbuild: *** compilation phase failed
```

Obviously, we can correct this error by putting the Add_One function at the same level as the Access To Function type, i.e. at library level:

Listing 177: access to subprogram types.ads

```
package Access_To_Subprogram_Types is
1
2
      type Access To Procedure is
3
        access procedure (I : in out Integer);
4
5
      type Access To Function is
6
        access function (I : Integer) return Integer;
7
8
9
```

```
end Access_To_Subprogram_Types;
```

Listing 178: add one.ads

```
function Add_One (I : Integer) return Integer;
```

Listing 179: add one.adb

```
function Add_One (I : Integer) return Integer is
1
  begin
2
      return I + 1;
3
  end Add_One;
4
```

Listing 180: show access to subprogram error.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_To_Subprogram_Types;
3
   use Access_To_Subprogram_Types;
4
5
   with Add One;
6
7
   procedure Show Access To Subprogram Error is
8
      Func : Access To Function;
9
10
      Value : Integer := 0;
11
   begin
12
      Func := Add_One'Access;
13
14
      Put Line ("Value: " & Value'Image);
15
      Value := Func (Value);
16
      Put Line ("Value: " & Value'Image);
17
   end Show Access To Subprogram Error;
18
```

Code block metadata

Project: Courses.Advanced Ada.Resource Management.Access Types.Accessibility Rules Access To Subprograms. Access To Subprogram Accessibility Error Less Deep Fix MD5: 7f7488c541fb457ced653a2e6cc2fad1

Runtime output

Value: 1	Value:	0
	Value:	1

As a recommendation, resolving accessibility issues in the case of access-to-subprograms is best done by refactoring the subprograms of your source code — for example, moving subprograms to a different level.

15.16.1 Unchecked Access

Previously, we discussed about the *Unchecked_Access attribute* (page 654), which we can use to circumvent accessibility issues in specific cases for access-to-objects. We also said in that section that this attribute only exists for objects, not for subprograms. We can use the previous example to illustrate this limitation:

Listing 181: access_to_subprogram_types.ads

```
package Access_To_Subprogram_Types is

type Access_To_Procedure is
access procedure (I : in out Integer);

type Access_To_Function is
access function (I : Integer) return Integer;
end Access_To_Subprogram_Types;
```

Listing 182: show_access_to_subprogram_error.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_To_Subprogram_Types;
3
   use Access_To_Subprogram_Types;
4
5
   procedure Show Access To Subprogram Error is
6
      Func : Access_To_Function;
7
8
      function Add One (I : Integer)
9
                return Integer is
10
         (I + 1);
11
12
      Value : Integer := 0;
13
   beain
14
      Func := Add One'Access;
15
16
      Put Line ("Value: " & Value'Image);
17
      Value := Func (Value);
18
      Put Line ("Value: " & Value'Image);
19
   end Show Access To Subprogram Error;
20
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Accessibility_Rules_ ⇔Access_To_Subprograms.Access_To_Subprogram_Accessibility_Error_Same_Lifetime MD5: clee1946f0c979eb30fbf2c72c426f50

Build output

```
gprbuild: *** compilation phase failed
```

When we analyze the Show_Access_To_Subprogram_Error procedure, we see that the Func object and the Add_One function have the same lifetime. Therefore, in this very specific case, we could safely assign Add_One'Access to Func and call Func for Value. Due to the accessibility rules, however, this assignment is illegal. (Obviously, the accessibility issue here is that the Access To Function type has a potentially longer lifetime.)

In the case of access-to-objects, we could use Unchecked_Access to enforce assignments that we consider safe after careful analysis. However, because this attribute isn't available for access-to-subprograms, the best solution is to move the subprogram to a level that allows the assignment to be legal, as we said before.

1 In the GNAT toolchain

1 2

3

4 5

6

7 8

9

GNAT offers an equivalent for Unchecked_Access that can be used for subprograms: the Unrestricted_Access attribute. Note, however, that this attribute is not portable.

Listing 183: access_to_subprogram_types.ads

```
package Access_To_Subprogram_Types is
```

```
type Access_To_Procedure is
access procedure (I : in out Integer);
```

```
type Access_To_Function is
  access function (I : Integer) return Integer;
```

```
end Access_To_Subprogram_Types;
```

Listing 184: show_access_to_subprogram_error.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_To_Subprogram_Types;
3
   use Access To Subprogram Types;
4
5
   procedure Show Access To Subprogram Error is
6
       Func : Access_To_Function;
7
8
       function Add One (I : Integer)
9
                return Integer is
10
         (I + 1);
11
12
      Value : Integer := 0;
13
   begin
14
      Func := Add_One'Unrestricted Access;
15
                        ~~~~<del>~</del>~~~~~
16
                Allowing access to local function
       - -
17
18
       Put Line ("Value: " & Value'Image);
19
       Value := Func (Value);
20
       Put_Line ("Value: " & Value'Image);
21
   end Show_Access_To_Subprogram_Error;
22
```

Code block metadata

D	
Runtime	ουτρυτ

Value: 0 Value: 1

As we can see, the Unrestricted_Access attribute can be safely used in this specific case to circumvent the accessibility rule limitation.

15.17 Access and Address

As we know, an access type is not a pointer, and it doesn't just indicate an address in memory. In fact, to represent an address in Ada, we use *the Address type* (page 123). Also, as we discussed earlier, we can use operators such as <, >, + and - for addresses. In contrast to that, those operators aren't available for access types — except, of course, for = and /=.

In certain situations, however, we might need to convert between access types and addresses. In this section, we discuss how to do so.

1 In the Ada Reference Manual

- 13.3 Operational and Representation Attributes²⁹¹
- 13.7 The Package System²⁹²

15.17.1 Address and access conversion

The generic System.Address_To_Access_Conversions package allows us to convert between access types and addresses. This might be useful for specific low-level operations. Let's see an example:

Listing 185: show address conversion.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with System.Address_To_Access_Conversions;
3
   with System.Address_Image;
4
5
   procedure Show_Address_Conversion is
6
7
8
      package Integer_AAC is
         new System.Address_To_Access_Conversions
9
          (Object => Integer);
10
      use Integer_AAC;
11
12
      subtype Integer_Access is
13
       Integer_AAC.Object_Pointer;
14
       -- This is similar to:
15
      - -
16
      -- type Integer Access is access all Integer;
17
18
      I : aliased Integer := 5;
19
20
      AI : Integer_Access := I'Access;
21
   begin
      Put_Line ("I'Address : "
22
```

(continues on next page)

²⁹¹ http://www.ada-auth.org/standards/22rm/html/RM-13-3.html
²⁹² http://www.ada-auth.org/standards/22rm/html/RM-13-7.html

```
& System.Address_Image (I'Address));
23
24
      Put_Line ("AI.all'Address : "
25
                 & System.Address_Image
26
                      (AI.all'Address));
27
28
      Put_Line ("To_Address (AI) : "
29
                 & System.Address_Image
30
                      (To_Address (AI)));
31
   end Show_Address_Conversion;
32
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Address. ⇔Address_Conversion MD5: 717532026247044a667b60f6c1e1c7da

Runtime output

I'Address : 00007FFD73B5D1F4 AI.all'Address : 00007FFD73B5D1F4 To Address (AI) : 00007FFD73B5D1F4

In this example, we instantiate the generic System.Address_To_Access_Conversions package using **Integer** as our target object type. This new package (Integer_AAC) has an Object_Pointer type, which is equivalent to a declaration such as **type Integer_Access is access all Integer**. (In this example, we declare Integer_Access as a subtype of Integer_AAC.Object_Pointer to illustrate that.)

The Integer_AAC package also includes the To_Address function, which converts an access object to an address. If the actual parameter is not null, To_Address returns the same information as if we were using the Address attribute for the designated object. In other words, To_Address (AI) = AI.all'Address when AI /= null.

If the access value is null, To_Address returns Null_Address, while .**all**'Address makes the *access check* (page 514) fail because we have to dereference the access object (via .**all**) before retrieving its address (via the **Address** attribute).

In addition to the To_Address function, the To_Pointer function is available to convert from an address to an object of access type. For example:

Listing 186: show address conversion.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
   with System;
                      use System;
2
3
   with System.Address_To_Access_Conversions;
4
   with System.Address Image;
5
6
   procedure Show_Address_Conversion is
7
8
      package Integer AAC is
9
         new System.Address_To_Access_Conversions
10
           (Object => Integer);
11
      use Integer AAC;
12
13
      subtype Integer Access is
14
        Integer_AAC.Object_Pointer;
15
16
                  : aliased Integer := 5;
      Т
17
      AI_1, AI_2 : Integer_Access;
18
```

```
А
                   : Address;
19
   begin
20
      AI_1 := I'Access;
21
      A := To_Address (AI_1);
22
      AI_2 := To_Pointer (A);
23
24
      Put_Line ("AI_1.all'Address : "
25
                 & System.Address_Image
26
                      (AI_1.all'Address));
27
      Put_Line ("AI_2.all'Address : "
28
                 & System.Address_Image
29
                      (AI_2.all'Address));
30
31
      if AI_1 = AI_2 then
32
          Put_Line ("AI_1 = AI_2");
33
       else
34
          Put_Line ("AI_1 /= AI_2");
35
       end if;
36
   end Show_Address_Conversion;
37
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Address.

⇔Address_Conversion

MD5: 5c6fc19ca1aa227feba97ea610dd9218
```

Runtime output

AI_1.all'Address : 00007FFC2485B18C
AI_2.all'Address : 00007FFC2485B18C
AI_1 = AI_2

Here, we convert the A address back to an access value by calling To_Pointer (A). (When running this object, we see that AI_1 and AI_2 have the same access value.)

Conversion of unbounded designated types

Note that the conversions might not work in all cases. For instance, when the designated type — indicated by the formal Object parameter of the generic Address_To_Access_Conversions package — is unbounded, the result of a call to To_Pointer may not have bounds.

Let's adapt the previous code example and replace the **Integer** type by the (unbounded) **String** type:

Listing 187: show address conversion.ad

```
with Ada.Text_IO; use Ada.Text_IO;
1
   with System;
                     use System;
2
3
   with System.Address_To_Access_Conversions;
4
   with System.Address_Image;
5
6
   procedure Show_Address_Conversion is
7
8
9
      package String_AAC is
         new System.Address_To_Access_Conversions
10
           (Object => String);
11
      use String_AAC;
12
13
```

```
subtype Integer_Access is
14
         String_AAC.Object_Pointer;
15
16
                   : aliased String := "Hello";
      S
17
      AI_1, AI_2 : Integer_Access;
18
                   : Address;
      А
19
   beain
20
      AI_1 := S'Access;
21
            := To_Address (AI_1);
      Α
22
23
      AI_2 := To_Pointer (A);
24
25
            WARNING: Result might not have bounds
26
27
      Put_Line ("AI_1.all'Address : "
28
                  & System.Address_Image
29
                      (AI_1.all'Address));
30
      Put_Line ("AI_2.all'Address : "
31
                  & System.Address_Image
32
                      (AI_2.all'Address));
33
34
       if AI_1 = AI_2 then
35
          Put_Line ("AI_1 = AI_2");
36
37
       else
          Put_Line ("AI_1 /= AI_2");
38
      end if;
39
40
      Put_Line ("AI_1: " & AI_1.all);
41
      Put_Line ("AI_2: " & AI_2.all);
42
43
            WARNING: As AI_2 might not have bounds
       - -
44
                      due to the call to To Pointer
45
                      the behavior of this call to
       - -
46
                      the "&" operator is
47
       - -
                      unpredictable.
48
   end Show_Address_Conversion;
49
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Access_Types.Access_Address.

⇔Address_Conversion

MD5: bladcaalf2cb4dfbd157aebf7893bd72
```

Build output

Runtime output

AI_1.all'Address : 00007FFD4958B858 AI_2.all'Address : 00007FFD4958B858 AI_1 = AI_2 AI_1: Hello AI_2: Hello

In this case, the call to To_Pointer (A) might not have bounds, so any operation on AI_2 might lead to unpredictable results.

1 In the Ada Reference Manual

• 13.7.2 The Package System.Address_To_Access_Conversions²⁹³

²⁹³ http://www.ada-auth.org/standards/22rm/html/RM-13-7-2.html

CHAPTER SIXTEEN

ANONYMOUS ACCESS TYPES

16.1 Named and Anonymous Access Types

The previous chapter dealt with access type declarations such as this one:

```
type Integer_Access is access all Integer;
```

procedure Add_One (A : Integer_Access);

In addition to named access type declarations such as the one in this example, Ada also supports anonymous access types, which, as the name implies, don't have an actual type declaration.

To declare an access object of anonymous type, we just specify the subtype of the object or subprogram we want to have access to. For example:

procedure Add_One (A : access Integer);

When we compare this example with the previous one, we see that the declaration A : Integer_Access becomes A : access Integer. Here, access Integer is the anonymous access type declaration, and A is an access object of this anonymous type.

To be more precise, A : **access Integer** is an *access parameter* (page 735) and it's specifying an *anonymous access-to-object type* (page 715). Another flavor of anonymous access types are *anonymous access-to-subprograms* (page 758). We discuss all these topics in more details later.

Let's see a complete example:

Listing 1: show_anonymous_access_types.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Anonymous Access Types is
3
      I Var : aliased Integer;
4
5
             : access Integer;
6
      Α
                ` Anonymous access type
      - -
7
   beain
8
      A := I Var'Access;
9
            ^ Assignment to object of
10
      - -
              anonymous access type.
       - -
11
12
      A.all := 22;
13
14
      Put Line ("A.all: " & Integer'Image (A.all));
15
   end Show Anonymous Access Types;
16
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_Types.Simple_Anonymous_Access_Types MD5: f0c92c76d970089c1d503c599d6869dd

Runtime output

A.all: 22

Here, A is an access object whose value is initialized with the access to I_Var. Because the declaration of A includes the declaration of an anonymous access type, we don't declare an extra Integer_Access type, as we did in previous code examples.

In the Ada Reference Manual

3.10 Access Types²⁹⁴

16.1.1 Relation to named types

Anonymous access types were not part of the first version of the Ada standard, which only had support for named access types. They were introduced later to cover some use-cases that were difficult — or even impossible — with access types.

In this sense, anonymous access types aren't just access types without names. Certain accessibility rules for anonymous access types are a bit less strict. In those cases, it might be interesting to consider using them instead of named access types.

In general, however, we should only use anonymous access types in those specific cases where using named access types becomes too cumbersome. As a general recommendation, we should give preference to named access types whenever possible. (Anonymous access-to-object types have *drawbacks that we discuss later* (page 718).)

16.1.2 Benefits of anonymous access types

One of the main benefits of anonymous access types is their flexibility: since there isn't an explicit access type declaration associated with them, we only have to worry about the subtype S we intend to access.

Also, as long as the subtype S in a declaration **access** S is always the same, no conversion is needed between two access objects of that anonymous type, and the S'Access attribute always works.

Let's see an example:

Listing 2: show.adb

```
with Ada.Text_IO; use Ada.Text_IO;
procedure Show (Name : String;
V : access Integer) is
begin
Put_Line (Name & ".all: "
V & Integer'Image (V.all));
end Show:
```

²⁹⁴ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

```
Listing 3: show_anonymous_access_types.adb
```

```
with Show;
1
2
   procedure Show Anonymous Access Types is
3
       I_Var : aliased Integer;
4
       Α
             : access Integer;
5
       В
             : access Integer;
6
   begin
7
       A := I_Var'Access;
8
       B := A;
9
10
       A.all := 22;
11
12
       Show ("A", A);
Show ("B", B);
13
14
   end Show Anonymous Access Types;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Types.Anonymous_Access_Object_Assignment

MD5: 2822ca0bd6ac251dccc1ced60747fbe1
```

Runtime output

A.all: 22 B.all: 22

In this example, we have two access objects A and B. Since they're objects of anonymous access types that refer to the same subtype **Integer**, we can assign A to B without a type conversion, and pass those access objects as an argument to the Show procedure.

(Note that the use of an access parameter in the Show procedure is for demonstration purpose only: a simply **Integer** as the type of this input parameter would have been more than sufficient to implement the procedure. Actually, in this case, avoiding the access parameter would be the recommended approach in terms of clean Ada software design.)

In contrast, if we had used named type declarations, the code would be more complicated and more limited:

Listing 4: aux.ads

```
package Aux is
type Integer_Access is access all Integer;
procedure Show (Name : String;
V : Integer_Access);
end Aux;
```

Listing 5: aux.adb

```
with Ada.Text_I0; use Ada.Text_I0;
package body Aux is
procedure Show (Name : String;
V : Integer_Access) is
begin
```

```
8 Put_Line (Name & ".all: "
9 & Integer'Image (V.all));
10 end Show;
11
12 end Aux;
```

Listing 6: show_anonymous_access_types.adb

```
with Aux; use Aux;
1
2
   procedure Show_Anonymous_Access_Types is
3
      -- I Var : aliased Integer;
4
5
      A : Integer Access;
6
      B : Integer Access;
7
   begin
8
      -- A := I Var'Access;
9
       - -
              ^ ERROR: non-local pointer cannot
10
                          point to local object.
      - -
11
12
      A := new Integer;
13
      B := A;
14
15
      A.all := 22;
16
17
      Show ("A", A);
18
      Show ("B", B);
19
   end Show Anonymous Access Types;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Types.Anonymous_Access_Object_Assignment

MD5: 681c2cf7f5e8d520490cc5594484ce69
```

Runtime output

A.all: 22 B.all: 22

Here, apart from the access type declaration (Integer_Access), we had to make two adaptations to convert the previous code example:

- 1. We had to move the Show procedure to a package (which we simply called Aux) because of the access type declaration.
- Also, we had to allocate an object for A instead of retrieving the access attribute of I_Var because we cannot use a pointer to a local object in the assignment to a nonlocal pointer, as indicate in the comments.

This restriction regarding non-local pointer assignments is an example of the stricter accessibility rules that apply to named access types. As mentioned earlier, the S'Access attribute always works when we use anonymous access types — this is not always the case for named access types.

Important

As mentioned earlier, if we want to use two access objects in an operation, the rule says that the subtype S of the anonymous type used in their corresponding declaration must match. In the following example, we can see how this rule

```
works:
                   Listing 7: show anonymous access subtype error.adb
   procedure Show Anonymous Access Subtype Error is
1
      subtype Integer_1_10 is Integer range 1 .. 10;
2
3
      I Var : aliased Integer;
4
      A : access Integer := I_Var'Access;
5
      В
            : access Integer 1 10;
6
   begin
7
      A := I_Var'Access;
8
9
      B := A;
10
      -- ^ ERROR: subtype doesn't match!
11
12
      B := I_Var'Access;
13
      -- ^ ERROR: subtype doesn't match!
14
   end Show Anonymous Access Subtype Error;
15
        Code block metadata
        Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_
         →Types.Anonymous_Access_Types.Anonymous_Access_Subtype_Error
        MD5: cecfe703ea8b42bad61c45f33cbcb67b
        Build output
   show_anonymous_access_subtype_error.adb:10:09: error: target_
     designated subtype not compatible with type "Standard.Integer"
   show_anonymous_access_subtype_error.adb:13:09: error: object subtype_
    →must statically match designated subtype
   gprbuild: *** compilation phase failed
   Even though Integer 1 10 is a subtype of Integer, we cannot assign A to B because the
   subtype that their access type declarations refer to - Integer and Integer 1 10, re-
   spectively — doesn't match. The same issue occurs when retrieving the access attribute
```

of I Var in the assignment to B.

The later sections on *anonymous access-to-object type* (page 715) and *anonymous access-to-subprograms* (page 758) cover more specific details on anonymous access types.

16.2 Anonymous Access-To-Object Types

In the *previous chapter* (page 593), we introduced named access-to-object types and used those types throughout the chapter. Also, in the *previous section* (page 711), we've seen some simple examples of anonymous access-to-object types:

In addition to parameters and objects, we can use anonymous access types in discriminants, components of array and record types, renamings and function return types. (We discuss *anonymous access discriminants* (page 725) and *anonymous access parameters* (page 735) later on.) Let's see a code example that includes all these cases:

Listing 8: all_anonymous_access_to_object_types.ads

```
package All_Anonymous_Access_To_Object_Types is
1
2
      procedure Add_One (A : access Integer) is null;
3
                               ^ Anonymous access type
4
       - -
5
      AI : access Integer;
6
           ^ Anonymous access type
7
       - -
8
      type Rec (AI : access Integer) is private;
9
                       ^ Anonymous access type
10
11
      type Access_Array is
12
         array (Positive range <>) of
13
           access Integer;
14
          ^ Anonymous access type
15
16
      Arr : array (1 .. 5) of access Integer;
17
                                ^ Anonymous access type
18
19
      AI_Renaming : access Integer renames AI;
20
                     ^ Anonymous access type
21
22
      function Init_Access_Integer
23
       return access Integer is (null);
24
               ^ Anonymous access type
25
26
   private
27
28
      type Rec (AI : access Integer) is record
29
                      ^ Anonymous access type
30
          Internal_AI : access Integer;
31
                         ^ Anonymous access type
32
33
      end record;
34
35
   end All_Anonymous_Access_To_Object_Types;
36
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_To_Object_Types.All_Anonymous_Access_To_Object_Types

MD5: 6533b22a4e4526702320cb327bf6f69a
```

In this example, we see multiple examples of anonymous access-to-object types:

- as the A parameter of the Add_One procedure;
- in the declaration of the AI access object;
- as the AI discriminant of the Rec type;
- as the component type of the Access_Array type;
- as the component type of the Arr array;
- in the AI_Renaming renaming;
- as the return type of the Init_Access_Integer;
- as the Internal_AI of component of the Rec type.

In the Ada Reference Manual

3.10 Access Types²⁹⁵

16.2.1 Not Null Anonymous Access-To-Object Types

As expected, **null** is a valid value for an anonymous access type. However, we can forbid **null** as a valid value by using **not null** in the anonymous access type declaration. For example:

```
Listing 9: all_anonymous_access_to_object_types.ads
```

```
package All_Anonymous_Access_To_Object_Types is
1
2
       procedure Add_One (A : not null access Integer)
3
        is null;
4
                               ^ Anonymous access type
5
6
       I : aliased Integer;
7
8
       AI : not null access Integer := I'Access;
9
            ^ Anonymous access type
       - -
10
                                          ~ ~ ~ ~ ~ ~ ~ ~ ~
       - -
11
       - -
                        Initialization required!
12
13
       type Rec (AI : not null access Integer) is
14
          private;
15
                       ^ Anonymous access type
16
17
       type Access_Array is
18
          array (Positive range <>) of
19
            not null access Integer;
20
            ^ Anonymous access type
21
22
       Arr : array (1 .. 5) of
23
        not null access Integer :=
24
       -- ^ Anonymous access type
25
           (others => I'Access);
26
            ^^^^
27
       - -
                 Initialization required!
28
29
       AI Renaming : not null access Integer
30
        renames AI;
31
                      ^ Anonymous access type
32
       - -
33
       function Init_Access_Integer
34
       return not null access Integer is (I'Access);
35
              ^ Anonymous access type
36
       - -
                                               ~~~~~~
       - -
37
                             Initialization required!
38
       - -
39
   private
40
41
       type Rec (AI : not null access Integer) is
42
       record
43
                       ^ Anonymous access type
44
          Internal AI : not null access Integer;
45
                         ^ Anonymous access type
46
47
```

²⁹⁵ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

48 end record;

49 50

end All_Anonymous_Access_To_Object_Types;

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_To_Object_Types.All_Not_Null_Anonymous_Access_To_Object_Types

MD5: 027430aa9d5e19979206110f5e260d13
```

As you might have noticed, we took the previous code example and used **not null** for each usage instance of the anonymous access type. In this sense, this version of the code example is very similar to the previous one. Note, however, that we now have to explicitly initialize some elements to avoid the Constraint_Error exception being raised at runtime. This is the case for example for the AI access object:

AI : not null access Integer := I'Access;

If we hadn't initialized AI explicitly with I'Access, it would have been set to **null**, which would fail the **not null** constraint of the anonymous access type. Similarly, we also have to initialize the Arr array and return a valid access object for the Init_Access_Integer function.

16.2.2 Drawbacks of Anonymous Access-To-Object Types

Anonymous access-to-object types have important drawbacks. For example, some features that are available for named access types aren't available for the anonymous access types. Also, most of the drawbacks are related to how anonymous access-to-object types can potentially make the allocation and deallocation quite complicated or even error-prone.

For starters, some pool-related features aren't available for anonymous access-to-object types. For example, we cannot specify which pool is going to be used in the allocation of an anonymous access-to-object. In fact, the memory pool selection is compiler-dependent, so we cannot rely on an object being allocated from a specific pool when using **new** with an anonymous access-to-object type. (In contrast, as we know, each named access type has an associated pool, so objects allocated via **new** will be allocated from that pool.) Also, we cannot identify which pool was selected for the allocation of a specific object, so we don't have any information to use for the deallocation of that object.

Because the pool selection is hidden from us, this makes the memory deallocation more complicated. For example, we cannot instantiate the Ada.Unchecked_Deallocation procedure for anonymous access types. Also, some of the methods we could use to circumvent this limitation are error-prone, as we discuss in this section.

Also, storage-related features aren't available: specifying the storage size — especially, specifying that the access type has a storage size of zero — isn't possible.

Missing features

Let's see a code example that shows some of the features that aren't available for anonymous access-to-object types:

Listing 10: missing_features.ads

```
with Ada.Unchecked_Deallocation;
package Missing_Features is
    -- We cannot specify which pool will be used
```

```
-- in the anonymous access-to-object
6
       -- allocation; the pool is selected by the
7
       -- compiler:
8
       IA : access Integer := new Integer;
9
10
11
          All the features below aren't available
12
       - -
          for an anonymous access-to-object:
       - -
13
14
15
       -- Having a specific storage pool associated
16
          with the access type:
17
18
       type String_Access is
        access String;
19
       -- Automatically creates
20
          String_Access'Storage_Pool
21
       - -
22
       type Integer_Access is
23
         access Integer
24
           with Storage Pool =>
25
                  String_Access'Storage Pool;
26
                            \^^^^<u>`</u>^^^
       - -
27
                  Using the pool from another
       - -
28
                  access type.
29
       - -
30
       -- Specifying a deallocation function for the
31
       - -
          access type:
32
       procedure Free is
33
         new Ada. Unchecked Deallocation
34
           (Object => Integer,
35
                  => Integer_Access);
            Name
36
37
       -- Specifying a limited storage size for
38
       -- the access type:
39
       type Integer_Access_Store_128 is
40
         access Integer
41
            with Storage_Size => 128;
42
43
       -- Limiting the storage size for the
44
       -- access type to zero:
45
       type Integer_Access_Store_0 is
46
          access Integer
47
            with Storage_Size => 0;
48
49
   end Missing Features;
50
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Object_Types.Missing_Anonymous_Access_To_Object_Features MD5: 87a5c1413a720da84fab414cf63236ec

In the Missing_Features package, we see some of the features that we cannot use for the anonymous **access Integer** type, but that are available for equivalent named access types:

- There's no specific memory pool associated with the access object IA. In contrast, named types — such as String_Access and Integer_Access — have an associated pool, and we can use the Storage_Pool aspect and the Storage_Pool attribute to customize them.
- We cannot instantiate the Ada.Unchecked_Deallocation procedure for the access

Integer type. However, we can instantiate it for named access types such as the Integer_Access type.

• We cannot use the Storage_Size attribute for the **access Integer** type, but we're allowed to use it with named access types, which we do in the declaration of the Integer_Access_Store_128 and Integer_Access_Store_0 types.

Dangerous memory deallocation

We might think that we could make up for the absence of the Ada. Unchecked_Deallocation procedure for anonymous access-to-object types by converting those access objects (of anonymous access types) to a named type that has the same designated subtype. For example, if we have an access object IA of an anonymous **access Integer** type, we can convert it to the named Integer_Access type, provided this named access type is compatible with the anonymous access type, e.g.:

```
type Integer_Access is access all Integer
```

Let's see a complete code example:

Listing 11: show_dangerous_deallocation.adb

```
with Ada.Unchecked Deallocation;
1
2
   procedure Show_Dangerous_Deallocation is
3
      type Integer_Access is
4
         access all Integer;
5
6
      procedure Free is
7
         new Ada.Unchecked Deallocation
8
           (Object => Integer,
9
            Name
                  => Integer Access);
10
11
      IA : access Integer;
12
   begin
13
      IA := new Integer;
14
      IA.all := 30;
15
16
      -- Potentially erroneous deallocation via type
17
          conversion:
18
      Free (Integer_Access (IA));
19
20
   end Show Dangerous Deallocation;
21
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Object_Types.Deallocation_Anonymous_Access_To_Object_Erronoeus MD5: 91e024a4338e2e4f8d5b308d95499c1c

This example declares the IA access object of the anonymous **access Integer** type. After allocating an object for IA via **new**, we try to deallocate it by first converting it to the Integer_Access type, so that we can call the Free procedure to actually deallocate the object. Although this code compiles, it'll only work if both **access Integer** and Integer_Access types are using the same memory pool. Since we cannot really determine this, the result is potentially erroneous: it'll work if the compiler selected the same pool, but it'll fail otherwise.

1 Important

Because allocating memory for anonymous access types is potentially dangerous, we can use the No_Anonymous_Allocators restriction — which is available since Ada 2012 — to prevent this kind of memory allocation being used in the code. For example:

```
Listing 12: show_dangerous_allocation.adb
```

```
pragma Restrictions (No_Anonymous_Allocators);
```

```
procedure Show_Dangerous_Allocation is
    IA : access Integer;
begin
    IA := new Integer;
    IA.all := 30;
end Show_Dangerous_Allocation;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Anonymous_Access_To_Object_Types.No_Anonymous_Allocators

MD5: 0976821ce632f9635e33fd4f79c81ecd
```

Build output

1 2 3

4

5

6

7

8

```
show_dangerous_allocation.adb:6:10: error: violation of restriction "No_

Anonymous_Allocators" at line 1

gprbuild: *** compilation phase failed
```

Possible solution using named access types

A better solution to avoid issues when allocating and deallocating memory for anonymous access-to-object types is to allocate the object using a known pool. As mentioned before, the memory pool associated with a named access type is well-defined, so we can use this kind of types for memory allocation. In fact, we can use a named memory type to allocate an object via **new**, and then associate this allocated object with the access object of anonymous access type.

Let's see a code example:

Listing 13: show_successful_deallocation.adb

```
with Ada.Unchecked Deallocation;
1
2
   procedure Show_Successful_Deallocation is
3
4
      type Integer Access is
5
6
         access Integer;
7
      procedure Free is
8
         new Ada.Unchecked Deallocation
9
           (Object => Integer,
10
            Name => Integer_Access);
11
12
                : access Integer;
      ΤA
13
      Typed_IA : Integer_Access;
14
15
   begin
16
      Typed IA := new Integer;
17
      IA := Typed_IA;
18
      IA.all := 30;
19
20
       -- Deallocation of the access object that has
21
      -- an associated type:
22
      Free (Typed_IA);
23
```

24

25 end Show_Successful_Deallocation;

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_To_Object_Types.Deallocation_Anonymous_Access_To_Object_1

MD5: eff8b54adfcc8cce10920dc3620ff1b9
```

In this example, all operations related to memory allocation are exclusively making use of the Integer_Access type, which is a named access type. In fact, **new Integer** allocates the object from the pool associated with the Integer_Access type, and the call to Free deallocates this object back into that pool. Therefore, associating this object with the IA access object — in the IA := Typed_IA assignment — doesn't creates problems afterwards in the object's deallocation. (When calling Free, we only refer to the object of named access type, so the object is deallocated from a known pool.)

Of course, a potential issue here is that IA becomes a *dangling reference* (page 652) after the call to Free. Therefore, we can improve this solution by completely hiding the memory allocation and deallocation for the anonymous access types in subprograms — e.g. as part of a package. By doing so, we don't expose the named access type, thereby reducing the possibility of dangling references.

In fact, we can generalize this approach with the following (generic) package:

Listing 14: hidden_anonymous_allocation.ads

```
generic
1
      type T is private;
2
   package Hidden Anonymous Allocation is
З
4
      function New T
5
         return not null access T;
6
7
      procedure Free (Obj : access T);
8
   end Hidden Anonymous Allocation;
10
```

Listing 15: hidden_anonymous_allocation.adb

```
with Ada.Unchecked Deallocation;
1
2
   package body Hidden_Anonymous_Allocation is
3
4
      type T_Access is access all T;
5
6
      procedure T Access Free is
7
         new Ada. Unchecked Deallocation
8
           (Object => T,
9
            Name
                   => T_Access);
10
11
      function New T
12
         return not null access T is
13
      begin
14
          return T_Access'(new T);
15
          -- Using allocation of the T_Access type:
16
              object is allocated from T_Access's pool
          - -
17
      end New T;
18
19
      procedure Free (Obj : access T) is
20
```

```
Tmp : T_Access := T_Access (Obj);
21
      begin
22
          T Access Free (Tmp);
23
          -- Using deallocation procedure of the
24
          - -
              T_Access type
25
       end Free;
26
27
   end Hidden_Anonymous_Allocation;
28
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Object_Types.Hidden_Alloc_Dealloc_Anonymous_Access_To_Object MD5: bd3831829f34f06a1d3c25a975c850a3

In the generic Hidden_Anonymous_Allocation package, New_T allocates a new object internally and returns an anonymous access to this object. The Free procedure deallocates this object.

In the body of the Hidden_Anonymous_Allocation package, we use the named access type T_Access to handle the actual memory allocation and deallocation. As expected, because those operations happen on the pool associated with the T_Access type, we don't have to worry about potential deallocation issues.

Finally, we can instantiate this package for the type we want to have anonymous access types for, say a type named Rec. Then, when using the Rec type in the main subprogram, we can simply call the corresponding subprograms for memory allocation and deallocation. For example:

```
Listing 16: info.ads
```

```
with Hidden Anonymous Allocation;
1
2
   package Info is
3
4
       type Rec is private;
5
6
       function New Rec return not null access Rec;
7
8
      procedure Free (Obj : access Rec);
9
10
   private
11
12
       type Rec is record
13
          I : Integer;
14
       end record;
15
16
       package Rec Allocation is new
17
         Hidden Anonymous Allocation (\top => Rec);
18
19
       function New Rec return not null access Rec
20
         renames Rec_Allocation.New_T;
21
22
       procedure Free (Obj : access Rec)
23
         renames Rec Allocation.Free;
24
25
   end Info;
26
```

Listing 17: show_info_allocation_deallocation.adb

```
with Info; use Info;
procedure Show_Info_Allocation_Deallocation is
RA : constant not null access Rec := New_Rec;
begin
Free (RA);
r end Show_Info_Allocation_Deallocation;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_To_Object_Types.Hidden_Alloc_Dealloc_Anonymous_Access_To_Object

MD5: d71e8ed70e280c6d5d9fc2d49c1eb6c3
```

In this example, we instantiate the Hidden_Anonymous_Allocation package in the Info package, which also defines the Rec type. We associate the New_T and Free subprograms with the Rec type by using subprogram renaming. Finally, in the Show_Info_Allocation_Deallocation procedure, we use these subprograms to allocate and deallocate the type.

Possible solution using the stack

Another approach that we could consider to avoid memory deallocation issues for anonymous access-to-object types is by simply using the stack for the object creation. For example:

Listing 18:	show	automatic	deallocation.adb

```
procedure Show Automatic Deallocation is
1
      I : aliased Integer;
2
           ^ Allocating object on the stack
3
4
      IA : access Integer;
5
   begin
6
      IA := I'Access;
7
      -- Indirect allocation:
8
      -- object creation on the stack.
9
10
      IA.all := 30;
11
12
      -- Automatic deallocation at the end of the
13
         procedure because the integer variable is
14
          on the stack.
15
   end Show_Automatic_Deallocation;
16
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_To_Object_Types.Deallocation_Anonymous_Access_To_Object_2

MD5: 4381db8ba87717978a9629b1e6a5f1fc
```

In this case, we create the I object on the stack by simply declaring it. Then, we get access to it and assign it to the IA access object.

With this approach, we're indirectly allocating an object for an anonymous access type by creating it on the stack. Also, because we know that the I is automatically deallocated when it gets out of scope, we don't have to worry about explicitly deallocating the object referred by IA.

When to use anonymous access-to-objects types

In summary, anonymous access-to-object types have many drawbacks that often outweigh *their benefits* (page 712). In fact, allocation for those types can quickly become very complicated. Therefore, in general, they're not a good alternative to named access types. Indeed, the difficulties that we've just seen might make them a much worse option than just using named access types instead.

We might consider using anonymous access-to-objects types only in cases when we reach a point in our implementation work where using named access types becomes impossible — or when using them becomes even more complicated than equivalent solutions using anonymous access types. This scenario, however, is usually the exception rather than the rule. Thus, as a general guideline, we should always aim to use named access types.

That being said, an important exception to this advice is when we're *interfacing to other languages* (page 738). In this case, as we'll discuss later, using anonymous access-to-objects types can be significantly simpler (compared to named access types) without the drawbacks that we've just discussed.

16.3 Access discriminants

Previously, we've discussed *discriminants as access values* (page 603). In that section, we only used named access types. Now, in this section, we see how to use anonymous access types as discriminants. This feature is also known as *access discriminants* and it provides some flexibility that can be interesting in terms of software design, as we'll discuss later.

Let's start with an example:

```
Listing 19: custom_recs.ads
```

```
package Custom_Recs is
1
2
      -- Declaring a discriminant with an anonymous
3
      -- access type:
4
      type Rec (IA : access Integer) is record
5
         I : Integer := IA.all;
6
      end record;
7
8
      procedure Show (R : Rec);
9
10
   end Custom_Recs;
11
```

Listing 20: custom_recs.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Custom Recs is
3
4
      procedure Show (R : Rec) is
5
      begin
6
          Put_Line ("R.IA = "
7
                    & Integer'Image (R.IA.all));
8
          Put_Line ("R.I = "
9
                    & Integer'Image (R.I));
10
11
      end Show;
12
   end Custom_Recs;
13
```

```
Listing 21: show_access_discriminants.adb
```

```
with Custom_Recs; use Custom_Recs;
1
2
   procedure Show Access Discriminants is
3
      I : aliased Integer := 10;
4
      R : Rec (I'Access);
5
   begin
6
      Show (R);
7
8
      I := 20;
9
      R.I := 30;
10
      Show (R);
11
   end Show_Access_Discriminants;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_

→Discriminants.Simple_Example

MD5: f8e127fda4f7ea0f1593165d6a966df6
```

Runtime output

```
R.IA = 10
R.I = 10
R.IA = 20
R.I = 30
```

In this example, we use an anonymous access type for the discriminant in the declaration of the Rec type of the Custom_Recs package. In the Show_Access_Discriminants procedure, we declare R and provide access to the local I integer.

Similarly, we can use unconstrained designated subtypes:

Listing 22: persons.ads

```
package Persons is
1
2
      -- Declaring a discriminant with an anonymous
3
      -- access type whose designated subtype is
4
5
         unconstrained:
      type Person (Name : access String) is record
6
         Age : Integer;
7
      end record;
8
9
      procedure Show (P : Person);
10
11
   end Persons;
12
```

Listing 23: persons.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Persons is
3
4
      procedure Show (P : Person) is
5
      begin
6
         Put Line ("Name = "
7
                    & P.Name.all);
8
         Put Line ("Age = "
9
                    & Integer'Image (P.Age));
10
```

```
end Show;
end Persons;
```

Listing 24: show_person.adb

```
with Persons; use Persons;
1
2
   procedure Show Person is
3
      S : aliased String := "John";
4
      P : Person (S'Access);
5
  begin
6
      P.Age := 30;
7
      Show (P):
8
  end Show Person;
a
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_

→Discriminants.Persons

MD5: f0149d572e0ec192476836bfdf00dd9e
```

Runtime output

Name = John Age = 30

In this example, for the discriminant of the Person type, we use an anonymous access type whose designated subtype is unconstrained. In the Show_Person procedure, we declare the P object and provide access to the S string.

1 In the Ada Reference Manual

- 3.7 Discriminants²⁹⁶
- 3.10.2 Operations of Access Types²⁹⁷

16.3.1 Default Value of Access Discriminants

In contrast to named access types, we cannot use a default value for the access discriminant of a non-limited type:

```
Listing 25: custom_recs.ads
```

```
package Custom_Recs is
1
2
      -- Declaring a discriminant with an anonymous
3
      -- access type and a default value:
4
      type Rec (IA : access Integer :=
5
                        new Integer'(0)) is
6
      record
7
         I : Integer := IA.all;
8
      end record;
9
10
  end Custom_Recs;
11
```

Code block metadata

²⁹⁶ http://www.ada-auth.org/standards/22rm/html/RM-3-7.html
 ²⁹⁷ http://www.ada-auth.org/standards/22rm/html/RM-3-10-2.html

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_ ⇔Discriminants.Default_Expression_Non_Limited_Type MD5: c3ddflcdfdaefa873ad66b9e47e03058

Build output

However, if we change the type declaration to be a limited type, having a default value for the access discriminant is OK:

Listing 26: custom recs.ads

```
package Custom Recs is
1
2
      -- Declaring a discriminant with an anonymous
3
      -- access type and a default value:
4
      type Rec (IA : access Integer :=
5
                        new Integer'(0)) is limited
6
      record
7
         I : Integer := IA.all;
8
      end record;
9
10
      procedure Show (R : Rec);
11
12
   end Custom Recs;
13
```

Listing 27: custom_recs.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Custom_Recs is
3
4
      procedure Show (R : Rec) is
5
      begin
6
          Put_Line ("R.IA = "
7
                    & Integer'Image (R.IA.all));
8
          Put_Line ("R.I = "
9
                    & Integer'Image (R.I));
10
      end Show;
11
12
   end Custom_Recs;
13
```

Code block metadata

Build output

```
custom_recs.ads:6:21: warning: coextension will not be deallocated when its

□associated owner is deallocated [enabled by default]
```

Note that, if we don't provide a value for the access discriminant when declaring an object R, the default value is allocated (via **new**) during R's creation.

Listing 28: show_access_discriminants.adb

```
with Custom Recs; use Custom Recs;
1
2
   procedure Show Access Discriminants is
3
      R : Rec;
4
5
      - -
          This triggers "new Integer'(0)", so an
      - -
6
          integer object is allocated and stored in
      - -
7
      - -
         the R.IA discriminant.
8
   begin
9
      Show (R);
10
11
      -- R gets out of scope here, and the object
12
      -- allocated via new hasn't been deallocated.
13
   end Show_Access_Discriminants;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_

→Discriminants.Default_Expression_Limited_Type

MD5: f5d9dee26044ccab2193ab419638de79
```

Build output

Runtime output

 $\begin{array}{rcl} \mathsf{R.IA} &= & \mathsf{0} \\ \mathsf{R.I} &= & \mathsf{0} \end{array}$

In this case, the allocated object won't be deallocated when R gets out of scope!

16.3.2 Benefits of Access Discriminants

Access discriminants have the same benefits that we've already seen earlier while discussing *discriminants as access values* (page 603). An additional benefit is its extended flexibility: access discriminants are compatible with any access T'Access, as long as T is of the designated subtype.

Consider the following example using the named access type Access_String:

```
Listing 29: persons.ads
```

```
package Persons is
1
2
      type Access_String is access all String;
3
4
       -- Declaring a discriminant with a named
5
       -- access type:
6
      type Person (Name : Access String) is record
7
         Age : Integer;
8
      end record;
9
10
      procedure Show (P : Person);
11
12
   end Persons;
13
```

```
Listing 30: persons.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Persons is
3
4
      procedure Show (P : Person) is
5
      begin
6
          Put_Line ("Name = "
7
                     & P.Name.all);
8
          Put_Line ("Age = "
9
                    & Integer'Image (P.Age));
10
      end Show;
11
12
   end Persons;
13
```

```
Listing 31: show_person.adb
```

```
with Persons; use Persons;
1
2
   procedure Show Person is
3
      S : aliased String := "John";
4
      P : Person (S'Access);
5
                    ^^^^ ERROR: cannot use local
       - -
6
       - -
                                     object
7
       - -
8
9
       -- We can, however, allocate the string via
10
      -- new:
11
       - -
       -- S : Access_String := new String'("John");
12
       - -
          P : Person (S);
13
   begin
14
      P.Age := 30;
15
      Show (P);
16
   end Show_Person;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_

⇔Discriminants.Persons

MD5: e918db3790c7ffeeb7c0f54ced9f48b9
```

Build output

```
show_person.adb:5:16: error: non-local pointer cannot point to local object
gprbuild: *** compilation phase failed
```

This code doesn't compile because we cannot have a non-local pointer (Access_String) pointing to the local object S. The only way to make this work is by allocating the string via **new** (i.e.: S : Access_String := **new String**).

However, if we use an access discriminant in the declaration of Person, the code compiles fine:

Listing 32: persons.ads

```
1 package Persons is
2
3 -- Declaring a discriminant with an anonymous
4 -- access type:
5 type Person (Name : access String) is record
```

```
6 Age : Integer;
7 end record;
8 
9 procedure Show (P : Person);
10 
11 end Persons;
```

```
Listing 33: show_person.adb
```

```
with Persons; use Persons;
1
2
   procedure Show Person is
3
      S : aliased String := "John";
4
      P : Person (S'Access);
5
                   ~~~~~ OK
6
7
      -- Allocating the string via new and using it
8
      -- in P's declaration is OK as well, but we
9
      -- should manually deallocate it before S
10
      -- gets out of scope:
11
12
      - -
      -- S : access String := new String'("John");
13
         P : Person (S);
      - -
14
  begin
15
      P.Age := 30;
16
      Show (P);
17
  end Show Person;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_

→Discriminants.Persons

MD5: 6516fb4e0cbbac9cfe07a56e48ea9ff3
```

Runtime output

Name = John Age = 30

In this case, getting access to the local object S and using it for P's discriminant is perfectly fine.

16.3.3 Preventing dangling pointers

Note that the usual rules that prevent dangling pointers still apply here. This ensures that we can safely use access discriminants. For example:

```
Listing 34: show_person.adb
```

```
with Persons; use Persons;
1
2
   procedure Show Person is
3
4
      function Local Init return Person is
5
         S : aliased String := "John";
6
      begin
7
         return (Name => S'Access, Age => 30);
8
                    ~~~~~~~~~~
9
         - -
                  ERROR: dangling reference!
         - -
10
      end Local Init;
11
```

```
12
13 P : Person := Local_Init;
14 begin
15 Show (P);
16 end Show_Person;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Access_

⇔Discriminants.Persons

MD5: 9c8d2aebf60b8bb19e455cb6bc5730eb
```

Build output

In this example, compilation fails in the Local_Init function when trying to return an object of Person type because S'Access would be a dangling reference.

16.4 Self-reference

Previously, we've seen that we can declare *self-references* (page 620) using named access types. We can do the same with anonymous access types. Let's revisit the code example that implements linked lists:

Listing	35:	linked	lists.ads

```
generic
1
       type T is private;
2
   package Linked_Lists is
3
4
       type List is limited private;
5
6
       procedure Append_Front
7
          (L : in out List;
8
           Ε:
                       T);
9
10
       procedure Append Rear
11
          (L : in out List;
12
           Ε:
                       T):
13
14
       procedure Show (L : List);
15
16
   private
17
18
       type Component is record
19
          Next : access Component;
20
                   ~~~~~~
          - -
21
          - -
                   Self-reference
22
          - -
23
                    (Note that we haven't finished the
          - -
24
                    declaration of the "Component" type
          - -
25
                   yet, but we're already referring to
26
          - -
                   it.)
27
          - -
28
          Value : T:
29
       end record;
30
```

```
31
32
32
33
34
end Linked_Lists;
```

Listing 36: linked_lists.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Linked_Lists is
3
4
       procedure Append_Front
5
          (L : in out List;
6
           Ε:
                       T)
7
       is
8
          New First : constant List := new
9
             Component'(Value => E,
10
                        Next => L);
11
       begin
12
          L := New_First;
13
       end Append Front;
14
15
       procedure Append_Rear
16
          (L : in out List;
17
           E :
                       T)
18
       is
19
          New Last : constant List := new
20
            Component'(Value => E,
21
22
                        Next => null);
       begin
23
          if L = null then
24
             L := New_Last;
25
          else
26
             declare
27
                 Last : List := L;
28
             begin
29
                 while Last.Next /= null loop
30
                    Last := List (Last.Next);
31
                     - -
32
                     - -
                          type conversion:
33
                             "access Component" to
                     - -
34
                             "List"
35
                    - -
                 end loop;
36
                 Last.Next := New_Last;
37
             end;
38
          end if;
39
       end Append_Rear;
40
41
       procedure Show (L : List) is
42
          Curr : List := L;
43
       begin
44
          if L = null then
45
              Put_Line ("[ ]");
46
          else
47
              Put ("[");
48
              loop
49
                 Put (Curr.Value'Image);
50
                 Put (" ");
51
                 exit when Curr.Next = null;
52
                 Curr := Curr.Next;
53
              end loop;
54
```

```
55 Put_Line ("]");
56 end if;
57 end Show;
58
59 end Linked_Lists;
```

Listing 37: test_linked_list.adb

```
with Linked Lists;
1
2
   procedure Test Linked List is
3
        package Integer Lists is new
4
          Linked Lists (T => Integer);
5
        use Integer_Lists;
6
7
        L : List;
8
   begin
9
        Append Front (L, 3);
10
        Append_Rear (L, 4);
11
        Append_Rear (L, 5);
12
        Append_Front (L, 2);
13
        Append_Front (L, 1);
14
        Append_Rear (L, 6);
15
        Append_Rear (L, 7);
16
17
        Show (L);
18
   end Test Linked List;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Self_

⊶Reference.Linked_List_Example

MD5: 98b9b2ce6fac3064326e6345520dc650
```

Runtime output

[1234567]

Here, in the declaration of the Component type (in the private part of the generic Linked_Lists package), we declare Next as an anonymous access type that refers to the Component type. (Note that at this point, we haven't finished the declaration of the Component type yet, but we're already using it as the designated subtype of an anonymous access type.) Then, we declare List as a general access type (with Component as the designated subtype).

It's worth mentioning that the List type and the anonymous **access** Component type aren't the same type, although they share the same designated subtype. Therefore, in the implementation of the Append_Rear procedure, we have to use type conversion to convert from the anonymous **access** Component type to the (named) List type.

16.5 Mutually dependent types using anonymous access types

In the section on *mutually dependent types using access types* (page 623), we've seen a code example that was using named access types. We could now rewrite it using anonymous access types:

Listing 38: mutually_dependent.ads

```
package Mutually Dependent is
1
2
      type T2;
3
4
      type T1 is record
5
          B : access T2;
6
      end record:
7
8
      type T2 is record
9
         A : access T1;
10
      end record;
11
12
   end Mutually_Dependent;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Mutually_

→Dependent_Anonymous_Access_Types.Example

MD5: 09f869d99b9c16882554588bb806a113
```

In this example, T1 and T2 are mutually dependent types. We're using anonymous access types in the declaration of the B and A components.

16.6 Access parameters

In the previous chapter, we talked about *parameters as access values* (page 610). As you might have expected, we can also use anonymous access types as parameters of a subprogram. However, they're limited to be **in** parameters of a subprogram or return type of a function (also called the access result type):

Listing 39: names.ads

```
package Names is
1
2
      function Init (S1, S2 : String)
3
                    return access String;
4
                    ^^^^
5
6
      -- Anonymous access type as the access
7
      - -
         result type.
8
      procedure Show (N : access constant String);
9
                         ^^^^
10
         Anonymous access type as a parameter type.
11
      - -
12
  end Names;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.Names

MD5: 622a76c4b133ed2715f18c175694cbe2
```

In this example, we have a string as the access result type of the Init function, and another string as the access parameter of the Show procedure.

This is the complete code example:

Listing 40: names.ads

```
package Names is
1
2
       function Init (S1, S2 : String)
3
                       return access String;
4
5
      procedure Show (N : access constant String);
6
7
   private
8
9
       function Init (S1, S2 : String)
10
                       return access String is
11
         (new String'(S1 & "-" & S2));
12
13
   end Names;
14
```

Listing 41: names.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Names is
3
4
      procedure Show (N : access constant String) is
5
      begin
6
          Put_Line ("Name: " & N.all);
7
      end Show;
8
9
   end Names;
10
```

Listing 42: show_names.adb

```
with Names; use Names;
procedure Show Names is
N : access String := Init ("Lily", "Ann");
begin
Show (N);
r end Show Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.Names

MD5: 9fe629f29de2898f2b82d9146b22fd1a
```

Runtime output

Name: Lily-Ann

Note that we're not using the **in** parameter mode in the Show procedure above. Usually, this parameter mode can be omitted, as it is the default parameter mode — **procedure** P (I : Integer) is the same as **procedure** P (I : in Integer). However, in the case of the Show procedure, the **in** parameter mode isn't just optionally absent. In fact, for access parameters, the parameter mode is always implied as **in**, so writing it explicitly is actually forbidden. In other words, we can only write N : **access String** or N : **access constant String**, but we cannot write N : **in access String** or N : **in access constant String**.

For further reading...

When we discussed *parameters as access values* (page 610) in the previous chapter, we saw how we can simply use different parameter modes to write a program instead of using access types. Basically, to implement the same functionality, we just replaced the access types by selecting the correct parameter modes instead and used *simpler* data types.

Let's do the same exercise again, this time by adapting the previous code example with anonymous access types:

Listing 43: names.ads

```
package Names is
```

private

1 2

3

4 5

6 7

8 9

10

11

12 13

14

1 2

3 4

5

6

7

8 9

1 2

3

4

5

6

```
function Init (S1, S2 : String)
    return String is
    (S1 & "-" & S2);
```

end Names;

```
Listing 44: names.adb
```

```
with Ada.Text_I0; use Ada.Text_I0;
package body Names is
    procedure Show (N : String) is
    begin
        Put_Line ("Name: " & N);
end Show;
```

```
10 end Names;
```

Listing 45: show_names.adb

```
with Names; use Names;
procedure Show_Names is
    N : String := Init ("Lily", "Ann");
begin
    Show (N);
end Show Names;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Anonymous_Access_Parameters.Names_String

MD5: 643f193999ef8de9bcefb11d9bdd21d7
```

Runtime output

Name: Lily-Ann

Although we're using simple strings instead of access types in this version of the code example, we're still getting a similar behavior. However, there is a small, yet important difference in the way the string returned by Init is being allocated: while the previous implementation (which was using an access result type) was allocating the string on the heap, we're now allocating the string on the stack.

Later on, we talk about the accessibility rules in the case of access parameters (page 757).

In general, we should avoid access parameters whenever possible and simply use objects and parameter modes directly, as it makes the design simpler and less error-prone. One exception is when we're interfacing to other languages, especially C: this is our *next topic* (page 738). Another time when access parameters are vital is for inherited primitive operations for tagged types. We discuss this *later on* (page 741).

1 In the Ada Reference Manual

• 3.10 Access Types²⁹⁸

16.6.1 Interfacing To Other Languages

We can use access parameters to interface to other languages. This can be particularly useful when interfacing to C code that makes use of pointers. For example, let's assume we want to call the add_one function below in our Ada implementation:

Listing 46: operations_c.h

```
void add_one(int *p_i);
```

Listing 47: operations c.c

```
1 void add_one(int *p_i)
2 {
3     *p_i = *p_i + 1;
4 }
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.C_Interfacing

MD5: 3270f3b2415266a203a6f4c605c3831b
```

We could map the **int** * parameter of add one to **access Integer** in the Ada specification:

```
procedure Add_One (IA : access Integer)
with Import, Convention => C;
```

This is a complete code example:

Listing 48: operations.ads

```
package Operations is
procedure Add_One (IA : access Integer)
with Import, Convention => C;
end Operations;
```

```
<sup>298</sup> http://www.ada-auth.org/standards/22rm/html/RM-3-10.html
```

```
Listing 49: show operations.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Operations; use Operations;
3
4
   procedure Show Operations is
5
      I : aliased Integer := 42;
6
   begin
7
      Put Line (I'Image);
8
      Add_One (I'Access);
9
      Put_Line (I'Image);
10
   end Show_Operations;
11
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_Parameters.C_Interfacing MD5: 0219acdbd2dad69962875199ffdd930e

Once again, we can replace access parameters with simpler types by using the appropriate parameter mode. In this case, we could replace **access Integer** by **aliased in out Integer**. This is the modified version of the code:

Listing 50: operations.ads

```
package Operations is
procedure Add_One
(IA : aliased in out Integer)
with Import, Convention => C;
end Operations;
```

Listing 51: show_operations.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Operations; use Operations;
3
4
   procedure Show Operations is
5
      I : aliased Integer := 42;
6
   begin
7
      Put Line (I'Image);
8
      Add One (I);
9
      Put Line (I'Image);
10
   end Show Operations;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.C_Interfacing

MD5: 2c5a81b8d77f0fff8a73f7912be6b6fe
```

However, there are situations where aliased objects cannot be used. For example, suppose we want to allocate memory inside a C function. In this case, the pointer to that memory block must be mapped to an access type in Ada.

Let's extend the previous C code example and introduce the alloc_integer and dealloc_integer functions, which allocate and deallocate an integer value:

```
Listing 52: operations_c.h
```

```
int * alloc_integer();
```

```
void dealloc_integer(int *p_i);
```

```
5 void add_one(int *p_i);
```

Listing 53: operations_c.c

```
#include <stdlib.h>
1
2
   int * alloc_integer()
3
   {
4
        return malloc(sizeof(int));
5
6
   }
7
   void dealloc_integer(int *p_i)
8
9
   {
10
        free (p_i);
11
   }
12
   void add_one(int *p_i)
13
   {
14
15
        *p_i = *p_i + 1;
   }
16
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_Parameters.C_Interfacing MD5: ec6dea12d0a948489cce21b0cc0a1ad2

In this case, we really have to use access types to interface to these C functions. In fact, we need an access result type to interface to the alloc_integer() function, and an access parameter in the case of the dealloc_integer() function. This is the corresponding specification in Ada:

Listing 54: operations.ads

```
package Operations is
1
2
      function Alloc Integer return access Integer
3
        with Import, Convention => C;
4
5
      procedure Dealloc Integer (IA : access Integer)
6
        with Import, Convention => C;
7
8
      procedure Add One
9
        (IA : aliased in out Integer)
10
           with Import, Convention => C;
11
12
   end Operations;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.C_Interfacing

MD5: bcbc8a87037b64fc6469e67b928e6172
```

Note that we're still using an aliased integer type for the Add_One procedure, while we're using access types for the other two subprograms.

Finally, as expected, we can use this specification in a test application:

Listing 55: show operations.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Operations; use Operations;
3
4
   procedure Show_Operations is
5
      I : access Integer := Alloc Integer;
6
   begin
7
      I.all := 42:
8
      Put Line (I.all'Image);
9
10
      Add One (I.all);
11
      Put Line (I.all'Image);
12
13
      Dealloc Integer (I);
14
   end Show_Operations;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.C_Interfacing

MD5: b2b96a166926528bc44059b56e31fb55
```

In this application, we get a C pointer from the alloc_integer function and encapsulate it in an Ada access type, which we then assign to I. In the last line of the procedure, we call Dealloc_Integer and pass I to it, which deallocates the memory block indicated by the C pointer.

1 In the Ada Reference Manual

3.10 Access Types²⁹⁹

16.6.2 Inherited Primitive Operations For Tagged Types

In order to declare inherited primitive operations for tagged types that use access types, we need to use access parameters. The reason is that, to be a primitive operation for some tagged type — and hence inheritable — the subprogram must reference the tagged type name directly in the parameter profile. This means that a named access type won't suffice, because only the access type name would appear in the profile. For example:

```
Listing 56: inherited_primitives.ads
```

```
package Inherited_Primitives is
1
2
      type T is tagged private;
3
4
      type T_Access is access all T;
5
6
      procedure Proc (N : T_Access);
7
       -- Proc is not a primitive of type T.
8
9
      type T_Child is new T with private;
10
11
      type T_Child_Access is access all T_Child;
12
13
```

(continues on next page)

²⁹⁹ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

```
14 private
15
16 type T is tagged null record;
17
18 type T_Child is new T with null record;
19
20 end Inherited_Primitives;
```

Listing 57: inherited_primitives.adb

```
with Ada.Text_IO; use Ada.Text_IO;
package body Inherited_Primitives is
procedure Proc (N : T_Access) is null;
end Inherited Primitives;
```

Listing 58: show_inherited_primitives.adb

```
with Inherited Primitives;
1
   use Inherited Primitives;
2
3
   procedure Show_Inherited_Primitives is
4
                : T_Access
5
      0bj
                             := new T;
      Obj_Child : T_Child_Access := new T_Child;
6
   begin
7
      Proc (Obj);
8
      Proc (Obj_Child);
9
10
            ERROR: Proc is not inherited!
11
   end Show Inherited Primitives;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.Inherited_Primitives

MD5: 8235b21caa9f1f105f533d74d891adfe
```

Build output

In this example, Proc is not a primitive of type T because it's referring to type T_Access, not type T. This means that Proc isn't inherited when we derive the T_Child type. Therefore, when we call Proc ($0bj_Child$), a compilation error occurs because the compiler expects type T_Access — there's no Proc ($N : T_Child_Access$) that could be used here.

If we replace T_Access in the Proc procedure with an an access parameter (**access** T), the subprogram becomes a primitive of T:

Listing	59:	inherited	primitives	.ads

```
package Inherited_Primitives is
```

```
type T is tagged private;
```

```
4
      procedure Proc (N : access T);
5
      -- Proc is a primitive of type T.
6
7
      type T_Child is new T with private;
8
9
   private
10
11
      type T is tagged null record;
12
13
      type T_Child is new T with null record;
14
15
   end Inherited_Primitives;
16
```

Listing 60: inherited primitives.adb

```
package body Inherited_Primitives is
procedure Proc (N : access T) is null;
end Inherited_Primitives;
```

Listing 61: show_inherited_primitives.adb

```
with Inherited_Primitives;
1
   use Inherited_Primitives;
2
3
   procedure Show_Inherited_Primitives is
4
              : access T
                            := new T;
      0bj
5
      Obj_Child : access T_Child := new T_Child;
6
   begin
7
      Proc (Obj);
8
9
      Proc (Obj_Child);
      - -
10
            OK: Proc is inherited!
11
      - -
  end Show_Inherited_Primitives;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_Parameters.Inherited_Primitives

MD5: a7e9b8bc92e346758cc4ade43bb4b02d
```

Now, the child type T_Child (derived from the T) inherits the primitive operation Proc. This inherited operation has an access parameter designating the child type:

```
type T_Child is new T with private;
procedure Proc (N : access T_Child);
-- Implicitly inherited primitive operation
```

In the Ada Reference Manual

• 3.9.2 Dispatching Operations of Tagged Types³⁰⁰

³⁰⁰ http://www.ada-auth.org/standards/22rm/html/RM-3-9-2.html

16.7 User-Defined References

Implicit dereferencing (page 625) isn't limited to the contexts that Ada supports by default: we can also add implicit dereferencing to our own types by using the Implicit_Dereference aspect.

To do this, we have to declare:

- a reference type, where we use the Implicit_Dereference aspect to specify the reference discriminant, which is the record discriminant that will be dereferenced; and
- a reference object, which contains an access value that will be dereferenced.

Also, for the reference type, we have to:

- specify the reference discriminant as an access discriminant (page 725); and
- indicate the name of the reference discriminant when specifying the Implicit_Dereference aspect.

Let's see a simple example:

Listing 62: show_user_defined_reference.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show_User_Defined_Reference is
3
4
       type Id_Number is record
5
          Id : Positive;
6
       end record;
7
8
9
       -- Reference type:
10
11
       type Id_Ref (Ref : access Id Number) is
12
                      ^ reference discriminant
13
         null record
14
          with Implicit_Dereference => Ref;
15
         - -
16
                            name of the reference
         - -
17
                            discriminant
         - -
18
19
20
       -- Access value:
21
22
       I : constant access Id_Number :=
23
             new Id Number'(Id => 42);
24
25
26
       - -
       -- Reference object:
27
28
       - -
       R : Id_Ref (I);
29
   begin
30
       Put_Line ("ID: "
31
                  & Positive'Image (R.Id));
32
                                      ^ Equivalent to:
33
                                          R.Ref.Id
       - -
34
       - -
                                        or:
35
                                         R.Ref.all.Id
36
   end Show User Defined Reference;
37
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_

→Defined_References.Simple_User_Defined_References

MD5: 33eaa7e8e75b4eb56d64dcc17e2932aa
```

Runtime output

ID: 42

Here, we declare a simple record type (Id_Number) and a corresponding reference type (Id_Ref). Note that:

- the reference discriminant Ref has an access to the Id_Number type; and
- we indicate this reference discriminant in the Implicit_Dereference aspect.

Then, we declare an access value (the I constant) and use it for the Ref discriminant in the declaration of the reference object R.

Finally, we implicitly dereference R and access the Id component by simply writing R.Id — instead of the extended forms R.Ref.Id or R.Ref.all.Id.

```
1 Important
```

The extended form mentioned in the example that we just saw (R.Ref.**all**.Id) makes it clear that two steps happen when evaluating R.Id:

- First, R.Ref is implied from R because of the Implicit_Dereference aspect.
- Then, R.Ref is implicitly dereferenced to R.Ref.all.

After these two steps, we can access the actual object. (In our case, we can access the Id component.)

Note that we cannot use access types directly for the reference discriminant. For example, if we made the following change in the previous code example, it wouldn't compile:

```
type Id_Number_Access is access Id_Number;
-- Reference type:
type Id_Ref (Ref : Id_Number_Access) is
-- ^ ERROR: it must be
-- an access
-- discriminant!
null record
with Implicit_Dereference => Ref;
```

However, we could use other forms — such as **not null access** — in the reference discriminant:

```
-- Reference type:
type Id_Ref (Ref : not null access Id_Number) is
null record
with Implicit_Dereference => Ref;
```

In the Ada Reference Manual

• 4.1.5 User-Defined References³⁰¹

³⁰¹ http://www.ada-auth.org/standards/22rm/html/RM-4-1-5.html

16.7.1 Dereferencing of tagged types

Naturally, implicit dereferencing is also possible when calling primitives of a tagged type. For example, let's change the declaration of the Id_Number type from the previous code example and add a Show primitive.

```
Listing 63: info.ads
```

```
package Info is
1
      type Id_Number (Id : Positive) is
2
        tagged private;
3
4
      procedure Show (R : Id Number);
5
  private
6
      type Id Number (Id : Positive) is
7
        tagged null record;
8
  end Info;
9
```

Listing 64: info.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body Info is
3
4
      procedure Show (R : Id_Number) is
5
      begin
6
          Put_Line ("ID: " & Positive'Image (R.Id));
7
      end Show;
8
9
   end Info;
10
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_

⇔Defined_References.Dereferencing_Tagged_Types

MD5: 4de65094963450dc3a7505dbf93c2551
```

Then, let's declare a reference type and a reference object in the test application:

Listing 65: show_user_defined_reference.adb

```
with Info; use Info;
1
2
   procedure Show_User_Defined_Reference is
3
4
       -- Reference type:
5
      type Id Ref (Ref : access Id Number) is
6
        null record
7
          with Implicit_Dereference => Ref;
8
9
      -- Access value:
10
      I : constant access Id Number :=
11
             new Id Number (42);
12
13
       -- Reference object:
14
      R : Id Ref (I);
15
   begin
16
17
      R.Show;
18
      -- Equivalent to:
19
       -- R.Ref.all.Show;
20
21
```

22 end Show_User_Defined_Reference;

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_

→Defined_References.Dereferencing_Tagged_Types

MD5: 9c5dfc4f2b8e085efde9e61689243f70
```

Runtime output

ID: 42

Here, we can call the Show procedure by simply writing R.Show instead of R.Ref.all.Show.

16.7.2 Simple container

A typical application of user-defined references is to create cursors when iterating over a container. As an example, let's implement the National_Date_Info package to store the national day of a country:

Listing	66:	national	date	info.ads

```
package National Date Info is
1
2
       subtype Country Code is String (1 .. 3);
3
4
      type Time is record
5
          Year : Integer;
6
          Month : Positive range 1 .. 12;
7
          Day : Positive range 1 ... 31;
8
      end record;
9
10
      type National Date is tagged record
11
          Country : Country_Code;
12
          Date
                  : Time;
13
       end record;
14
15
       type National Date Access is
16
        access National Date;
17
18
      procedure Show (Nat Date : National Date);
19
20
   end National Date Info;
21
```

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body National Date Info is
3
4
      procedure Show (Nat Date : National Date) is
5
      begin
6
          Put_Line ("Country: "
7
                    & Nat_Date.Country);
8
          Put_Line ("Year:
9
                    & Integer'Image
10
                         (Nat_Date.Date.Year));
11
      end Show;
12
13
   end National_Date_Info;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_

→Defined_References.National_Dates

MD5: 90fd6740d701025e1d5f30c9751a528d
```

Here, National_Date is a record type that we use to store the national day information. We can call the Show procedure to display this information.

Now, let's implement the National_Date_Containers with a container for national days:

Listing 68: national_date_containers.ads

```
with National_Date_Info; use National_Date_Info;
1
2
   package National Date Containers is
3
4
      -- Reference type:
5
      type National Date Reference
6
        (Ref : access National_Date) is
7
          tagged limited null record
8
            with Implicit_Dereference => Ref;
9
10
      -- Container (as an array):
11
      type National_Dates is
12
        array (Positive range <>) of
13
          National_Date_Access;
14
15
         The Find function scans the container to
16
         find a specific country, which is returned
17
      - -
      -- as a reference object.
18
      function Find (Nat_Dates : National_Dates;
19
                      Country : Country_Code)
20
                      return National_Date_Reference;
21
22
   end National Date Containers;
23
```

Listing 69: national_date_containers.adb

```
package body National_Date_Containers is
1
2
      function Find (Nat_Dates : National_Dates;
3
                       Country : Country Code)
4
5
                       return National_Date_Reference
6
      is
7
      begin
          for I in Nat_Dates'Range loop
8
             if Nat_Dates (I).Country = Country then
9
                return National_Date_Reference'(
10
                          Ref => Nat_Dates (I));
11
                            ~~~~~
                - -
12
                - -
                     Returning reference object with a
13
                - -
                     reference to the national day we
14
                - -
                      found.
15
             end if;
16
17
         end loop;
18
19
          return
           National_Date_Reference'(Ref => null);
20
                            ~~~~~~~~~~~~~~~~~~~~~~
21
          - -
               Returning reference object with a null
22
               reference in case the country wasn't
23
```

```
24 -- found. This will trigger an exception
25 -- if we try to dereference it.
26 end Find;
27
28 end National_Date_Containers;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_

→Defined_References.National_Dates

MD5: ec37ae93a7052c4bc731b2a7be0763ab
```

Package National_Date_Containers contains the National_Dates type, which is an array type for declaring containers that we use to store the national day information. We can also see the declaration of the National_Date_Reference type, which is the reference type returned by the Find function when looking for a specific country in the container.

Important

We're declaring the container type (National_Dates) as an array type just to simplify the code. In many cases, however, this approach isn't recommended! Instead, we should use a private type in order to encapsulate — and better protect — the information stored in the actual container.

Finally, let's see a test application that stores information for some countries into the Nat_Dates container and displays the information for a specific country:

Listing	70:	show	national	dates.adb

```
with National Date Info;
1
   use National Date Info;
2
3
   with National Date Containers;
4
   use National Date Containers;
5
6
   procedure Show_National_Dates is
7
8
      Nat Dates : constant National Dates (1 .. 5) :=
9
         (new National_Date'("USA",
10
                              Time'(1776, 7, 4)),
11
          new National_Date'("FRA"
12
                              Time'(1789, 7, 14)),
13
          new National_Date'("DEU"
14
                              Time'(1990, 10, 3)),
15
          new National_Date'("SPA"
16
                              Time'(1492, 10, 12)),
17
          new National Date'("BRA",
18
                              Time'(1822, 9, 7)));
19
20
   begin
21
      Find (Nat Dates, "FRA").Show;
22
                               ^ implicit dereference
23
   end Show National Dates;
24
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_

⇔Defined_References.National_Dates

MD5: 771ecb91e8f890d4bb9b08115ae833f4
```

Runtime output

Country:	FRA			
Year:	1789			

Here, we call the Find function to retrieve a reference object, whose reference (access value) has the national day information of France. We then implicitly dereference it to get the tagged object (of National_Date type) and display its information by calling the Show procedure.

1 Relevant topics

The National_Date_Containers package was implemented specifically as an accompanying package for the National_Date_Info package. It is possible, however, to generalize it, so that we can reuse the container for other record types. In fact, this is actually very straightforward:

```
Listing 71: generic containers.ads
```

```
generic
1
       type T is private;
2
       type T_Access is access T;
3
       type T_Cmp is private;
4
      with function Matches (E
                                     : T_Access;
5
                                Elem : T_Cmp)
6
                                return Boolean;
7
   package Generic_Containers is
8
9
       type Ref_Type (Ref : access T) is
10
         tagged limited null record
11
           with Implicit_Dereference => Ref;
12
13
       type Container is
14
         array (Positive range <>) of
15
           T_Access;
16
17
       function Find (Cont : Container;
18
                       Elem : T_Cmp)
19
                       return Ref_Type;
20
```

end Generic_Containers;

```
Listing 72: generic_containers.adb
```

```
package body Generic_Containers is
1
2
       function Find (Cont : Container;
3
                       Elem : T Cmp)
4
                       return Ref_Type is
5
       begin
6
          for I in Cont'Range loop
7
             if Matches (Cont (I), Elem) then
8
                return Ref_Type'(Ref => Cont (I));
9
             end if;
10
11
          end loop;
12
          return Ref_Type'(Ref => null);
13
      end Find;
14
15
   end Generic_Containers;
16
```

Code block metadata

21

22

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.User_ →Defined_References.National_Dates MD5: 94c23a48131a47439b5b41e985c3d6c1

When comparing the **Generic**_Containers package to the National_Date_Containers package, we see that the main difference is the addition of the Matches function, which indicates whether the current element we're evaluating in the for-loop of the Find function is the one we're looking for.

In the main application, we can implement the Matches function and declare the National_Date_Containers package as an instance of the **Generic**_Containers package:

Listing 73: show national dates.adb

```
with Generic Containers;
with National_Date_Info; use National_Date_Info;
procedure Show_National_Dates is
   function Matches Country
         : National_Date_Access;
     (E
     Elem : Country_Code)
      return Boolean is
        (E.Country = Elem);
   package National Date Containers is new
     Generic Containers
       (T
                 => National Date,
        T_Access => National_Date_Access,
        T Cmp
               => Country Code,
       Matches => Matches Country);
   use National_Date_Containers;
   subtype National Dates is Container;
   Nat Dates : constant
                 National Dates (1 .. 5) :=
     (new National_Date'("USA",
                         Time'(1776, 7, 4)),
      new National Date'("FRA"
                         Time'(1789, 7, 14)),
      new National_Date'("DEU"
                         Time'(1990, 10, 3)),
      new National_Date'("SPA"
                         Time'(1492, 10, 12)),
      new National_Date'("BRA"
                         Time'(1822, 9, 7)));
begin
   Find (Nat Dates, "FRA").Show;
end Show_National_Dates;
Code block metadata
```

Runtime output

Country: FRA Year: 1789

1

2 3

> 4 5

6

7

8

9

10 11

12

13

14

15

16

17 18

19 20

21 22

23

24

25

26

27

28

29

30

31

32

33

34 35

36

37

38

Here, we instantiate the **Generic**_Containers package with the Matches_Country function, which is an expression function that compares the country component of the current National_Date reference with the name of the country we desire to learn about.

This generalized approach is actually used for the standard containers from the Ada. Containers packages. For example, the Ada.Containers.Vectors is specified as follows:

with Ada.Iterator_Interfaces;

```
generic
  type Index_Type is range <>;
   type Element_Type is private;
  with function "=" (Left, Right : Element_Type)
                      return Boolean is <>;
package Ada.Containers.Vectors
 with Preelaborate, Remote_Types,
       Nonblocking,
       Global => in out synchronized is
   -- OMITTED
   type Reference_Type
     (Element : not null access Element Type) is
       private
         with Implicit_Dereference => Element,
              Nonblocking,
              Global => in out synchronized,
              Default Initial Condition =>
                (raise Program Error);
   -- OMITTED
   function Reference
     (Container : aliased in out Vector;
     Index : in Index_Type)
      return Reference_Type
        with Pre => Index in
                       First_Index (Container) ..
                       Last Index (Container)
                     or else raise
                             Constraint Error,
           Post =>
             Tampering_With_Cursors_Prohibited
               (Container),
           Nonblocking,
           Global => null,
           Use_Formal => null;
   -- OMITTED
   function Reference
     (Container : aliased in out Vector;
      Position : in Cursor)
      return Reference Type
        with Pre => (Position /= No_Element
                      or else raise
                              Constraint_Error)
                      and then
                        (Has Element
                          (Container, Position)
                         or else raise
                                 Program Error),
```

```
Post =>
Tampering_With_Cursors_Prohibited
  (Container),
Nonblocking,
Global => null,
Use Formal => null;
```

-- OMITTED

end Ada.Containers.Vectors;

(Note that most parts of the Vectors package were omitted for clarity. Please refer to the Ada Reference Manual for the complete package specification.)

Here, we see that the Implicit_Dereference aspect is used in the declaration of **Reference_Type**, which is the reference type returned by the Reference functions for an index or a cursor.

Also, note that the Vectors package has a formal equality function (=) instead of the Matches function we were using in our **Generic**_Containers package. The purpose of the formal function, however, is basically the same.

1 In the Ada Reference Manual

A.18.2 The Generic Package Containers. Vectors³⁰²

16.8 Anonymous Access Types and Accessibility Rules

In general, the *accessibility rules* (page 647) we've seen earlier also apply to anonymous access types. However, there are some subtle differences, which we discuss in this section.

Let's adapt the *code example from that section* (page 647) to make use of anonymous access types:

Listing 74: library level.ads

```
package Library_Level is
L0_A0 : access Integer;
L0_Var : aliased Integer;
end Library_Level;
```

Listing 75: show_library_level.adb

```
with Library_Level; use Library_Level;
1
2
   procedure Show_Library_Level is
3
      L1_Var : aliased Integer;
4
5
      L1_A0 : access Integer;
6
7
      procedure Test is
8
         L2_A0 : access Integer;
9
10
         L2_Var : aliased Integer;
11
```

(continues on next page)

³⁰² http://www.ada-auth.org/standards/22rm/html/RM-A-18-2.html

12

(continued from previous page)

```
begin
          L1_A0 := L2_Var'Access;
13
14
           - -
                     ILLEGAL: L2 object to
           - -
15
                               L1 access object
           - -
16
17
          L2_A0 := L2_Var'Access;
18
                     ~~~~~
19
           - -
           - -
                     LEGAL: L2 object to
20
                             L2 access object
21
       end Test;
22
23
24
    begin
       L0_A0 := new Integer'(22);
25
                  ~~~~~
26
       - -
                 LEGAL: L0 object to
27
       - -
                         L0 access object
       - -
28
29
       L0_A0 := L1_Var'Access;
30
31
                 ILLEGAL: L1 object to
       - -
32
                            L0 access object
       - -
33
34
       L1_A0 := L0_Var'Access;
35
36
       - -
                 LEGAL: L0 object to
37
       - -
38
       - -
                         L1 access object
39
       L1_A0 := L1_Var'Access;
40
                  ~~
41
                 LEGAL: L1 object to
       - -
42
                         L1 access object
43
44
       L0_A0 := L1_A0; -- legal!!
45
46
       - -
                 LEGAL:
                           L1 access object to
47
       - -
       - -
                            L0 access object
48
       - -
49
                 ILLEGAL: L1 object
       - -
50
                            (L1_A0 = L1_Var'Access)
       - -
51
       - -
                            to
52
       - -
                            L0 access object
53
       - -
54
                 This is actually OK at compile time,
       - -
55
                 but the accessibility check fails at
56
       - -
       - -
                 runtime.
57
58
       Test;
59
   end Show_Library_Level;
60
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.
Accessibility_Levels_Rules_Introduction.Accessibility_Library_Level
MD5: 255bdecebdaa735408db082edd583a0c
```

Build output

show_library_level.adb:13:16: error: non-local pointer cannot point to local object show_library_level.adb:30:13: error: non-local pointer cannot point to local object gprbuild: *** compilation phase failed

As we see in the code, in general, most accessibility rules are the same as the ones we've discussed when using named access types. For example, an assignment such as $L0_A0 := L1_Var'Access$ is illegal because we're trying to assign to an access object of less deep level.

However, assignment such as $L0_A0 := L1_A0$ are possible now: we don't get a type mismatch — as we did with named access types — because both objects are of anonymous access types. Note that the accessibility level cannot be determined at compile time: L1_A0 can hold an access value at library level (which would make the assignment legal) or at a deeper level. Therefore, the compiler introduces an accessibility check here.

However, the accessibility check used in L0_A0 := L1_A0 fails at runtime because the corresponding access value (L1_Var'Access) is of a deeper level than L0_A0, which is illegal. (If you comment out the L1_A0 := L1_Var'Access assignment prior to the L0_A0 := L1_A0 assignment, this accessibility check doesn't fail anymore.)

16.8.1 Conversions between Anonymous and Named Access Types

In the previous sections, we've discussed accessibility rules for named and anonymous access types separately. In this section, we see that the same accessibility rules apply when mixing both flavors together and converting objects of anonymous to named access types.

Let's adapt parts of the previous *code example* (page 647) and add anonymous access types to it:

Listing	76:	library_	leve	l.ads
---------	-----	----------	------	-------

```
package Library Level is
1
2
      type LO Integer Access is
3
        access all Integer;
4
5
      L0 Var : aliased Integer;
6
7
      L0_IA : L0_Integer_Access;
8
      L0 A0 : access Integer;
9
10
   end Library_Level;
11
```

Listing 77: show library level.adb

```
with Library_Level; use Library_Level;
1
2
   procedure Show_Library_Level is
3
      type L1 Integer Access is
4
        access all Integer;
5
6
      L1_IA : L1_Integer_Access;
7
      L1 A0 : access Integer;
8
9
      L1_Var : aliased Integer;
10
11
   begin
12
13
       -- From named type to anonymous type
14
15
16
      L0_IA := new Integer'(22);
17
      L1_IA := new Integer'(42);
18
19
```

```
L0A0 := L0IA;
20
                ~~~/
       - -
21
                LEGAL: assignment from
       - -
22
                        L0 access object (named type)
23
       - -
24
       - -
                        to
                        L0 access object
25
       - -
                          (anonymous type)
26
       - -
27
       L0_A0 := L1_IA;
28
                ~~~
       - -
29
       - -
                ILLEGAL: assignment from
30
31
       - -
                          L1 access object (named type)
32
       - -
                          to
                          L0 access object
33
       - -
                            (anonymous type)
34
       - -
35
       L1_A0 := L0_IA;
36
                 ~~~
37
       - -
                LEGAL: assignment from
       - -
38
       - -
                        L0 access object (named type)
39
       - -
                        to
40
                        L1 access object
       - -
41
                          (anonymous type)
       - -
42
43
       L1_A0 := L1_IA;
44
                 ~~~
45
       - -
                LEGAL: assignment from
46
       - -
                        L1 access object (named type)
47
       - -
                        to
48
       - -
                        L1 access object
       - -
49
                          (anonymous type)
50
51
       52
       -- From anonymous type to named type
53
             54
55
       L0_A0 := L0_Var'Access;
56
       L1_A0 := L1_Var'Access;
57
58
       L0_IA := L0_Integer_Access (L0_A0);
59
       - -
60
                LEGAL: conversion / assignment from
       - -
61
       - -
                        L0 access object
62
                          (anonymous type)
       - -
63
                        to
       - -
64
                        L0 access object (named type)
       - -
65
66
       L0_IA := L0_Integer_Access (L1_A0);
67
                ~~~~~~~~~~~
68
       - -
                ILLEGAL: conversion / assignment from
       - -
69
       - -
                          L1 access object
70
       - -
                             (anonymous type)
71
       - -
                          to
72
                          L0 access object (named type)
73
       - -
                          (accessibility check fails)
74
       - -
75
       L1_IA := L1_Integer_Access (L0_A0);
76
77
       - -
                LEGAL: conversion / assignment from
       - -
78
                        L0 access object
       - -
79
                          (anonymous type)
80
       - -
```

```
to
81
                         L1 access object (named type)
82
83
       L1_IA := L1_Integer_Access (L1_A0);
84
85
       - -
                 LEGAL: conversion / assignment from
       - -
86
                         L1 access object
87
       - -
                           (anonymous type)
88
       - -
                         to
89
                         L1 access object (named type)
90
   end Show Library Level;
91
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Accessibility_Levels_Rules_Introduction.Accessibility_Named_Anonymous_Access_

⇔Type_Conversions

MD5: a2e73bb0ed543bc4973850c80f951039
```

Build output

As we can see in this code example, mixing access objects of named and anonymous access types doesn't change the accessibility rules. Again, the rules are only violated when the target object in the assignment is *less* deep. This is the case in the L0_A0 := L1_IA and the L0_IA := L0_Integer_Access (L1_A0) assignments. Otherwise, mixing those access objects doesn't impose additional hurdles.

16.8.2 Accessibility rules on access parameters

In the previous chapter, we saw that the accessibility rules also apply to *access values as subprogram parameters* (page 650). In the case of access parameters, the rules are a bit less strict (as you may generally expect for anonymous access types), and the accessibility rules are checked at runtime. This allows use to use access values that would be illegal in the case of named access types because of their accessibility levels.

Let's adapt a previous code example to make use of access parameters:

Listing 78: names.ads

```
1 package Names is
2
3 procedure Show (N : access constant String);
4
5 end Names;
```

Listing 79: names.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   -- with Ada.Characters.Handling;
3
      use Ada.Characters.Handling;
4
   - -
   package body Names is
6
7
      procedure Show (N : access constant String) is
8
      begin
9
                                                                         (continues on next page)
```

```
10 -- for I in N'Range loop
11 -- N (I) := To_Lower (N (I));
12 -- end loop;
13 Put_Line ("Name: " & N.all);
14 end Show;
15
16 end Names;
```

Listing 80: show_names.adb

```
with Names; use Names;
procedure Show_Names is
    S : aliased String := "John";
    begin
    Show (S'Access);
    end Show_Names;
```

Code block metadata

Runtime output

Name: John

As we've seen in the previous chapter, compilation fails when we use named access types in this code example. In the case of access parameters, using S'Access doesn't make the compilation fail, nor does the accessibility check fail at runtime because S is still in scope when we call the Show procedure.

16.9 Anonymous Access-To-Subprograms

In the previous chapter, we talked about *named access-to-subprogram types* (page 677). Now, we'll see that the anonymous version of those types isn't much different from the named version.

Let's start our discussion by declaring a subprogram parameter using an anonymous access-to-procedure type:

```
Listing 81: anonymous_access_to_subprogram.ads
```

```
1 package Anonymous_Access_To_Subprogram is
2
3 procedure Proc
4 (P : access procedure (I : in out Integer));
5
6 end Anonymous Access To Subprogram;
```

Listing 82: anonymous_access_to_subprogram.adb

```
1 package body Anonymous_Access_To_Subprogram is
2
3 procedure Proc
4 (P : access procedure (I : in out Integer))
5 is
```

```
6 I : Integer := 0;
7 begin
8 P (I);
9 end Proc;
10
11 end Anonymous_Access_To_Subprogram;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Subprograms.Anonymous_Access_To_Subprogram_Example MD5: 2cbe76d7e23905d575bd27e29d5e3175

In this example, we use the anonymous **access procedure** (I : in out Integer) type as a parameter of the Proc procedure. Note that we need an identifier in the declaration: we cannot leave I out and write **access procedure** (in out Integer).

Before we look at a test application that makes use of the Anonymous_Access_To_Subprogram package, let's implement two simple procedures that we'll use later on:

Listing 83: add ten.ads

```
procedure Add_Ten (I : in out Integer);
```

Listing 84: add_ten.adb

```
procedure Add_Ten (I : in out Integer) is
begin
I := I + 10;
end Add Ten;
```

Listing 85: add_twenty.ads

```
procedure Add_Twenty (I : in out Integer);
```

Listing 86: add_twenty.adb

```
1 procedure Add_Twenty (I : in out Integer) is
2 begin
3 I := I + 20;
4 end Add_Twenty;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_
⇔Access_To_Subprograms.Anonymous_Access_To_Subprogram_Example
MD5: 50eaeaf27caaa9618b35ecdf8acc11fe
```

Finally, this is our test application:

Listing 87: show anonymous access to subprograms.adb

```
with Anonymous_Access_To_Subprogram;
use Anonymous_Access_To_Subprogram;
with Add_Ten;
procedure Show_Anonymous_Access_To_Subprograms is
begin
(continues on next page)
```

8	<pre>Proc (Add_Ten'Access);</pre>
9	^ Getting access to Add_Ten
10	procedure and passing it
11	to Proc
12	<pre>end Show_Anonymous_Access_To_Subprograms;</pre>

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Subprograms.Anonymous_Access_To_Subprogram_Example MD5: 13143ccf9620d26031484ba160a58fe1

Here, we get access to the Add_Ten procedure and pass it to the Proc procedure. Note that this implementation is not different from the *example for named access-to-subprogram types* (page 679). In fact, in terms of usage, anonymous access-to-subprogram types are very similar to named access-to-subprogram types. The major differences can be found in the corresponding *accessibility rules* (page 767).

1 In the Ada Reference Manual

• 3.10 Access Types³⁰³

16.9.1 Examples of anonymous access-to-subprogram usage

In the section about *named access-to-subprogram types* (page 677), we've seen a couple of different usages for those types. In all those examples we discussed, we could instead have used anonymous access-to-subprogram types. Let's see a code example that illustrates that:

Listing 88: all_anonymous_access_to_subprogram.ads

```
package All_Anonymous_Access_To_Subprogram is
1
2
3
          Anonymous access-to-subprogram as
4
       - -
       - -
          subprogram parameter:
5
6
      procedure Proc
7
       (P : access procedure (I : in out Integer));
8
9
10
       -- Anonymous access-to-subprogram in
11
       -- array type declaration:
12
13
      type Access_To Procedure Array is
14
        array (Positive range <>) of
15
           access procedure (I : in out Integer);
16
17
      protected type Protected Integer is
18
19
         procedure Mult Ten;
20
21
         procedure Mult_Twenty;
22
23
      private
24
        I : Integer := 1;
25
```

(continues on next page)

³⁰³ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

```
end Protected_Integer;
26
27
28
          Anonymous access-to-subprogram as
       - -
29
          component of a record type.
30
       - -
31
       type Rec_Access_To_Procedure is record
32
          AP : access procedure (I : in out Integer);
33
       end record;
34
35
36
       - -
          Anonymous access-to-subprogram as
37
38
       - -
           discriminant:
39
       type Rec_Access_To_Procedure_Discriminant
40
              (AP : access procedure
41
                       (I : in out Integer)) is
42
       record
43
          I : Integer := 0;
44
       end record;
45
46
       procedure Process
47
         (R : in out
48
                 Rec_Access_To_Procedure_Discriminant);
49
50
51
       generic
52
          type T is private;
53
54
          - -
              Anonymous access-to-subprogram as
          - -
55
              formal parameter:
          - -
56
57
          Proc_T : access procedure
58
                       (Element : in out T);
59
       procedure Gen_Process (Element : in out T);
60
61
   end All_Anonymous_Access_To_Subprogram;
62
```

Listing 89: all_anonymous_access_to_subprogram.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   package body All_Anonymous_Access_To_Subprogram is
3
4
      procedure Proc
5
         (P : access procedure (I : in out Integer))
6
      is
7
          I : Integer := 0;
8
      begin
9
          Put Line
10
            ("Calling procedure for Proc...");
11
          P (I);
12
          Put_Line ("Finished.");
13
      end Proc;
14
15
      procedure Process
16
         (R : in out
17
                Rec_Access_To_Procedure_Discriminant)
18
      is
19
      begin
20
          Put_Line
21
```

```
("Calling procedure for"
22
             & " Rec_Access_To_Procedure_Discriminant"
23
             & " type...");
24
          R.AP(R.I);
25
          Put_Line ("Finished.");
26
       end Process;
27
28
       procedure Gen_Process (Element : in out T) is
29
       begin
30
          Put Line
31
            ("Calling procedure for Gen Process...");
32
          Proc_T (Element);
33
          Put_Line ("Finished.");
34
       end Gen_Process;
35
36
       protected body Protected_Integer is
37
38
          procedure Mult_Ten is
39
          begin
40
             I := I * 10;
41
          end Mult_Ten;
42
43
          procedure Mult_Twenty is
44
          begin
45
              I := I * 20;
46
          end Mult_Twenty;
47
48
       end Protected_Integer;
49
50
   end All_Anonymous_Access_To_Subprogram;
51
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_

⇔Access_To_Subprograms.Anonymous_Access_To_Subprogram_Example

MD5: 628dcfdc5fe9b712f33fa044057093c2
```

In the All_Anonymous_Access_To_Subprogram package, we see examples of anonymous access-to-subprogram types:

- as a subprogram parameter;
- in an array type declaration;
- as a component of a record type;
- · as a record type discriminant;
- · as a formal parameter of a generic procedure.

Let's implement a test application that makes use of this package:

Listing 90: show_anonymous_access_to_subprograms.adb

```
with Ada.Text_IO; use Ada.Text_IO;
with Add_Ten;
with Add_Twenty;
with Add_Twenty;
with All_Anonymous_Access_To_Subprogram;
procedure Show_Anonymous_Access_To_Subprograms is
```

```
- -
          Anonymous access-to-subprogram as
11
       - -
           an object:
12
13
       Ρ
           : access procedure (I : in out Integer);
14
15
16
       - -
          Array of anonymous access-to-subprogram
       - -
17
       - -
           components
18
19
       PA : constant
20
              Access_To_Procedure_Array (1 .. 2) :=
21
                 (Add_Ten'Access,
22
                  Add_Twenty'Access);
23
24
25
       - -
       -- Anonymous array of anonymous
26
          access-to-subprogram components:
       - -
27
28
       PAA : constant
29
              array (1 .. 2) of access
30
                 procedure (I : in out Integer) :=
31
                   (Add_Ten'Access,
32
                    Add_Twenty'Access);
33
34
35
       - -
36
       -- Record with anonymous
       -- access-to-subprogram components:
37
38
       RA : constant Rec_Access_To_Procedure :=
39
               (AP => Add_Ten'Access);
40
41
42
       -- Record with anonymous
43
          access-to-subprogram discriminant:
       - -
44
45
       RD : Rec_Access_To_Procedure_Discriminant
46
               (AP => Add_Twenty'Access) :=
47
                 (AP => Add_Twenty'Access, I => 0);
48
49
50
       -- Generic procedure with formal anonymous
51
          access-to-subprogram:
       - -
52
53
       procedure Process Integer is new
54
         Gen_Process (T
                          => Integer,
55
                       Proc_T => Add_Twenty'Access);
56
57
58
       - -
          Object (APP) of anonymous
       - -
59
       - -
           access-to-protected-subprogram:
60
61
       PI : Protected Integer;
62
       APP : constant access protected procedure :=
63
               PI Mult Ten Access;
64
65
       Some_Int : Integer := 0;
66
67
   begin
       Put_Line ("Some_Int: " & Some_Int'Image);
68
69
70
```

10

```
-- Using object of
71
       -- anonymous access-to-subprogram type:
72
73
       P := Add_Ten'Access;
74
       Proc (P);
75
       P (Some_Int);
76
77
       P := Add_Twenty'Access;
78
       Proc (P);
79
       P (Some_Int);
80
81
       Put_Line ("Some_Int: " & Some_Int'Image);
82
83
84
       -- Using array with component of
85
           anonymous access-to-subprogram type:
86
       - -
87
        Put Line
88
           ("Calling procedure from PA array...");
89
90
       for I in PA'Range loop
91
          PA (I) (Some_Int);
92
          Put_Line ("Some_Int: " & Some_Int'Image);
93
       end loop;
94
95
       Put_Line ("Finished.");
96
97
       Put_Line
98
         ("Calling procedure from PAA array...");
99
100
       for I in PA'Range loop
101
           PAA (I) (Some_Int);
102
           Put_Line ("Some_Int: " & Some_Int'Image);
103
       end loop;
104
105
       Put_Line ("Finished.");
106
107
       Put_Line ("Some_Int: " & Some_Int'Image);
108
109
110
       -- Using record with component of
111
       -- anonymous access-to-subprogram type:
112
113
       RA.AP (Some_Int);
114
       Put Line ("Some Int: " & Some Int'Image);
115
116
117
       -- Using record with discriminant of
118
           anonymous access-to-subprogram type:
       - -
119
120
       Process (RD);
121
       Put_Line ("RD.I: " & RD.I'Image);
122
123
124
           Using procedure instantiated with
125
       - -
           formal anonymous access-to-subprogram:
       - -
126
127
       Process Integer (Some Int);
128
       Put_Line ("Some_Int: " & Some_Int'Image);
129
130
131
```

```
132 -- Using object of anonymous
133 -- access-to-protected-subprogram type:
134 --
135 APP.all;
136 end Show_Anonymous_Access_To_Subprograms;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Subprograms.Anonymous_Access_To_Subprogram_Example MD5: ec770c17e880a98fd2e9ab0110d4a858

Runtime output

```
Some Int: 0
Calling procedure for Proc...
Finished.
Calling procedure for Proc...
Finished.
Some Int:
          30
Calling procedure from PA array...
Some Int: 40
Some Int:
          60
Finished.
Calling procedure from PAA array...
Some Int: 70
Some Int: 90
Finished.
Some Int:
          90
Some Int: 100
Calling procedure for Rec Access To Procedure Discriminant type...
Finished.
RD.I: 20
Calling procedure for Gen Process...
Finished.
Some Int: 120
```

In the Show_Anonymous_Access_To_Subprograms procedure, we see examples of anonymous access-to-subprogram types in:

- in objects (P) and (APP);
- in arrays (PA and PAA);
- in records (RA and RD);
- in the binding to a formal parameter (Proc_T) of an instantiated procedure (Process_Integer);
- as a parameter of a procedure (Proc).

Because we already discussed all these usages in the section about *named access-to-subprogram types* (page 677), we won't repeat this discussion here. If anything in this code example is still unclear to you, make sure to revisit that section from the previous chapter.

16.9.2 Application of anonymous access-to-subprogram types

In general, there isn't much that speaks against using anonymous access-to-subprogram types. We can say, for example, that they're much more useful than *anonymous access-to-objects types* (page 715), which have *many drawbacks* (page 718) — as we discussed earlier.

There isn't much to be concerned when using anonymous access-to-subprogram types. For example, we cannot allocate or deallocate a subprogram. As a consequence, we won't have storage management issues affecting these types because the access to those subprograms will always be available and no memory leak can occur.

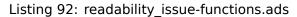
Also, anonymous access-to-subprogram types can be easier to use than named access-tosubprogram types because of their less strict *accessibility rules* (page 767). Some of the accessibility issues we might encounter when using named access-to-subprogram types can be solved by declaring them as anonymous types. (We discuss the accessibility rules of anonymous access-to-subprogram types in the next section.)

16.9.3 Readability

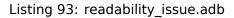
Note that readability suffers if you use a *cascade* of anonymous access-to-subprograms. For example:

```
Listing 91: readability issue.ads
```

```
package Readability_Issue is
1
2
       function F
3
         return access
4
           function (A : Integer)
5
                      return access
6
                        function (B : Float)
7
                                   return Integer;
8
9
   end Readability Issue;
10
```



```
package Readability_Issue.Functions is
1
2
      function To_Integer (V : Float)
3
                             return Integer is
4
         (Integer (V));
5
6
      function Select_Conversion
7
         (A : Integer)
8
          return access
9
            function (B : Float)
10
                       return Integer is
11
         (To_Integer'Access);
12
13
   end Readability_Issue.Functions;
14
```



```
with Readability Issue.Functions;
1
   use Readability_Issue.Functions;
2
3
   package body Readability Issue is
4
5
      function F
6
         return access
7
           function (A : Integer)
8
                      return access
9
                        function (B : Float)
10
                            return Integer is
11
         (Select Conversion'Access);
12
13
```

14 end Readability_Issue;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Subprograms.Readability_Issue MD5: 9e2ac58942c97b44c0d847c28e39bd11

In this example, the definition of F might compile fine, but it's simply too long to be readable. Not only that: we need to carry this *chain* to other functions as well — such as the Select_Conversion function above. Also, using these functions in an application is not straightforward:

```
Listing 94: show readability issue.adb
```

```
with Readability Issue;
1
   use Readability Issue;
2
   procedure Show Readability Issue is
4
      F1 : access
5
              function (A : Integer)
6
                        return access
7
                          function (B : Float)
8
                                     return Integer
9
            := F;
10
      F2 : access function (B : Float)
11
                              return Integer
12
            := F1(2);
13
      I : Integer := F2 (0.1);
14
15
   begin
      I := F1 (2) (0.1);
16
   end Show Readability Issue;
17
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.Anonymous_ ⇔Access_To_Subprograms.Readability_Issue MD5: 80267b1d673663e3cacba0c4978e6abf

Therefore, our recommendation is to avoid this kind of *access cascading* by carefully designing your application. In general, you won't need that.

16.10 Accessibility Rules and Anonymous Access-To-Subprograms

In principle, the *accessibility rules for anonymous access types* (page 753) that we've seen before apply to anonymous access-to-subprograms as well. Also, we had a discussion about *accessibility rules and access-to-subprograms* (page 702) in the previous chapter. In this section, we review some of the rules that we already know and discuss how they relate to anonymous access-to-subprograms.

1 In the Ada Reference Manual

3.10 Access Types³⁰⁴

³⁰⁴ http://www.ada-auth.org/standards/22rm/html/RM-3-10.html

16.10.1 Named vs. anonymous access-to-subprograms

Let's see an example of a named access-to-subprogram type:

```
Listing 95: show_access_to_subprogram_error.adb
```

```
with Ada.Text IO; use Ada.Text IO;
1
2
   procedure Show Access To Subprogram Error is
3
4
       type PI is access
5
         procedure (I : in out Integer);
6
7
      P : PI;
8
9
      I : Integer := 0;
10
   begin
11
       declare
12
          procedure Add One (I : in out Integer) is
13
          begin
14
             I := I + 1;
15
          end Add One;
16
      begin
17
          P := Add_One'Access;
18
       end:
19
   end Show_Access_To_Subprogram_Error;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Accessibility_Rules_Anonymous_Access_To_Subprograms.Simple_Example_Named

MD5: 41c36426112e799210b7704dd43b6217
```

Build output

In this example, we get a compilation error because the lifetime of the Add_One procedure is shorter than the access type PI.

In contrast, using an anonymous access-to-subprogram type eliminates the compilation error, i.e. the assignment P := Add One'Access becomes legal:

```
Listing 96: show_access_to_subprogram_error.adb
```

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   procedure Show Access To Subprogram Error is
3
      P : access procedure (I : in out Integer);
4
5
      I : Integer := 0;
6
   begin
7
      declare
8
          procedure Add One (I : in out Integer) is
9
          begin
10
             I := I + 1;
11
         end Add_One;
12
      begin
13
         P := Add_One'Access;
14
          -- RUNTIME ERROR: Add_One is out-of-scope
15
                              after this line.
16
```

17 end;

18 end Show_Access_To_Subprogram_Error;

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types. ⇔Accessibility_Rules_Anonymous_Access_To_Subprograms.Simple_Example_Anonymous MD5: a5eeb4a716b4f6a932dd74c580a07b66

Runtime output

raised PROGRAM_ERROR : show_access_to_subprogram_error.adb:14 accessibility check ⇔failed

In this case, the compiler introduces an accessibility check, which fails at runtime because the lifetime of Add_One is shorter than the lifetime of the access object P.

16.10.2 Named vs. anonymous access-to-subprograms as parameters

Using anonymous access-to-subprograms as parameters allows us to pass subprograms at any level. For certain applications, the restrictions that are applied to named access types might be too strict, so using anonymous access-to-subprograms might be a good way to circumvent those restrictions. They also allow the component developer to be independent of the clients' specific access types.

Note that the increased flexibility for anonymous access-to-subprograms means that some of the checks that are performed at compile time for named access-to-subprograms are done at runtime for anonymous access-to-subprograms.

Named access-to-subprograms as a parameter

Let's see an example using a named access-to-procedure type:

Listing 97: access_to_subprogram_types.ads

```
package Access_To_Subprogram_Types is
1
2
      type Integer_Array is
3
        array (Positive range <>) of Integer;
4
5
      type Process_Procedure is
6
        access
7
           procedure (Arr : in out Integer_Array);
8
9
      procedure Process
10
         (Arr : in out Integer_Array;
11
         Ρ
                       Process Procedure);
12
             1
13
   end Access_To_Subprogram_Types;
14
```

Listing 98: access to subprogram types.adb

```
1 package body Access_To_Subprogram_Types is
2
3 procedure Process
4 (Arr : in out Integer_Array;
5 P : Process Procedure) is
```

```
6 begin
7 P (Arr);
8 end Process;
9
10 end Access_To_Subprogram_Types;
```

Listing 99: show_access_to_subprogram_error.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Access_To_Subprogram_Types;
3
   use Access_To_Subprogram_Types;
4
5
   procedure Show_Access_To_Subprogram_Error is
6
7
       procedure Add One
8
         (Arr : in out Integer Array) is
9
       begin
10
          for E of Arr loop
11
             E := E + 1;
12
          end loop;
13
       end Add_One;
14
15
       procedure Display
16
        (Arr : in out Integer Array) is
17
       begin
18
          for I in Arr'Range loop
19
             Put Line ("Arr (" &
20
                        Integer'Image (I)
21
                        & "): "
22
                       & Integer'Image (Arr (I)));
23
          end loop;
24
       end Display;
25
26
       Arr : Integer_Array (1 .. 3) := (1, 2, 3);
27
   begin
28
       Process (Arr, Display'Access);
29
30
       Put Line ("Add One...");
31
       Process (Arr, Add_One'Access);
32
33
       Process (Arr, Display'Access);
34
   end Show_Access_To_Subprogram_Error;
35
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Accessibility_Rules_Anonymous_Access_To_Subprograms.Access_To_Subprogram_

⇔Parameter_Named

MD5: 76b70b52a0374fe0fd398024fe869876
```

Build output

In this example, we declare the Process_Procedure type in the Access_To_Subprogram_Types package and use it in the Process procedure, which we call in the Show_Access_To_Subprogram_Error procedure. The accessibility rules trigger a compilation error because the accesses (Add_One'Access and Display'Access) are at a deeper level than the access-to-procedure type (Process_Procedure).

As we know already, there's no Unchecked_Access attribute that we could use here. An easy way to make this code compile could be to move Add_One and Display to the library level.

Anonymous access-to-subprograms as a parameter

To circumvent the compilation error, we could also use anonymous access-to-subprograms instead:

Listing 100: access_to_subprogram_types.ads

```
package Access To Subprogram Types is
1
2
      type Integer_Array is
3
        array (Positive range <>) of Integer;
4
5
      procedure Process
6
        (Arr : in out Integer_Array;
7
         Ρ
            : access procedure
8
                  (Arr : in out Integer Array));
9
10
   end Access To Subprogram Types;
11
```

Listing 101: access_to_subprogram_types.adb

```
package body Access To Subprogram Types is
1
2
      procedure Process
3
4
         (Arr : in out Integer_Array;
         P : access procedure
5
                  (Arr : in out Integer_Array)) is
6
      begin
7
         P (Arr);
8
      end Process;
9
10
   end Access_To_Subprogram_Types;
11
```

Listing 102: show_access_to_subprogram_error.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Access To Subprogram Types;
3
   use Access To Subprogram Types;
4
   procedure Show_Access_To_Subprogram_Error is
6
7
      procedure Add One
8
        (Arr : in out Integer_Array) is
9
      begin
10
         for E of Arr loop
11
            E := E + 1;
12
         end loop;
13
      end Add One;
14
15
      procedure Display
16
```

```
(Arr : in out Integer_Array) is
17
      begin
18
          for I in Arr'Range loop
19
             Put_Line ("Arr (" &
20
                        Integer'Image (I)
21
                        & "): "
22
                       & Integer'Image (Arr (I)));
23
          end loop;
24
      end Display;
25
26
      Arr : Integer_Array (1 .. 3) := (1, 2, 3);
27
   beain
28
      Process (Arr, Display'Access);
29
30
      Put_Line ("Add_One...");
31
      Process (Arr, Add_One'Access);
32
33
      Process (Arr, Display'Access);
34
   end Show_Access_To_Subprogram_Error;
35
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⊲Accessibility_Rules_Anonymous_Access_To_Subprograms.Access_To_Subprogram_

⊲Parameter_Anonymous

MD5: a500e0a864f0adadc1d6823c1f50bd64
```

Runtime output

Arr (1): 1 Arr (2): 2 Arr (3): 3 Add_One... Arr (1): 2 Arr (2): 3 Arr (3): 4

Now, the code is accepted by the compiler because anonymous access-to-subprograms used as parameters allow passing of subprograms at any level. Also, we don't see a run-time exception because the subprograms are still *accessible* when we call Process.

16.10.3 Iterator

A typical example that illustrates well the necessity of using anonymous access-tosubprograms is that of a container iterator. In fact, many of the standard Ada containers the child packages of Ada.Containers — make use of anonymous access-to-subprograms for their Iterate subprograms.

1 In the Ada Reference Manual

- A.18.2 The Package Containers. Vectors³⁰⁵
- A.18.4 Maps³⁰⁶
- A.18.7 Sets³⁰⁷

³⁰⁵ http://www.ada-auth.org/standards/22rm/html/RM-A-18-2.html

³⁰⁶ http://www.ada-auth.org/standards/22rm/html/RM-A-18-4.html

³⁰⁷ http://www.ada-auth.org/standards/22rm/html/RM-A-18-7.html

Using named access-to-subprograms

Let's start with a simplified container type (Data_Container) using a named access-tosubprogram type (Process_Element) for iteration:

```
Listing 103: data_processing.ads
```

```
generic
1
      type Element is private;
2
   package Data Processing is
3
4
      type Data Container (Last : Positive) is
5
        private;
6
7
      Data_Container_Full : exception;
8
9
      procedure Append (D : in out Data_Container;
10
                          Ε:
                                      Element);
11
12
      type Process Element is
13
        not null access procedure (E : Element);
14
15
      procedure Iterate
16
         (D : Data_Container;
17
         Proc : Process_Element);
18
19
   private
20
21
      type Data_Container_Storage is
22
         array (Positive range <>) of Element;
23
24
      type Data_Container (Last : Positive) is
25
       record
26
          S
               : Data Container Storage (1 .. Last);
27
         Curr : Natural := 0;
28
      end record;
29
30
   end Data_Processing;
31
```

Listing 104: data_processing.adb

```
package body Data_Processing is
1
2
      procedure Append (D : in out Data_Container;
3
                          E: Element) is
4
      begin
5
          if D.Curr < D.S'Last then</pre>
6
             D.Curr := D.Curr + 1;
7
             D.S (D.Curr) := E;
8
          else
9
             raise Data_Container_Full;
10
             -- NOTE: This is just a dummy
11
                       implementation. A better
             - -
12
                       strategy is to add actual error
             - -
13
                       handling when the container is
             - -
14
                       full.
15
          end if;
16
      end Append;
17
18
      procedure Iterate
19
20
         (D : Data_Container;
         Proc : Process_Element) is
21
```

```
22 begin
23 for I in D.S'First .. D.Curr loop
24 Proc (D.S (I));
25 end loop;
26 end Iterate;
27
28 end Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Accessibility_Rules_Anonymous_Access_To_Subprograms.Iterator_Named

MD5: e48e8200e571b62d027753ee96c47fcb
```

In this example, we declare the Process_Element type in the generic Data_Processing package, and we use it in the Iterate procedure. We then instantiate this package as Float_Data_Processing, and we use it in the Show_Access_To_Subprograms procedure:

Listing 105: float_data_processing.ads

```
with Data_Processing;
package Float_Data_Processing is
new Data_Processing (Element => Float);
```

Listing 106: show_access_to_subprograms.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Float_Data_Processing;
3
   use Float_Data_Processing;
4
5
   procedure Show_Access_To_Subprograms is
6
7
      procedure Display (F : Float) is
8
      begin
9
         Put Line ("F :" & Float'Image (F));
10
      end Display;
11
12
      D : Data_Container (5);
13
14
   begin
        Append (D, 1.0);
15
        Append (D, 2.0);
16
        Append (D, 3.0);
17
18
        Iterate (D, Display'Access);
19
   end Show_Access_To_Subprograms;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⊶Accessibility_Rules_Anonymous_Access_To_Subprograms.Iterator_Named

MD5: 64ee435aac5f2817b7d9cecf538a1e4c
```

Build output

Using Display'Access in the call to Iterate triggers a compilation error because its life-

time is shorter than the lifetime of the Process_Element type.

Using anonymous access-to-subprograms

Now, let's use an anonymous access-to-subprogram type in the Iterate procedure:

```
Listing 107: data_processing.ads
```

```
generic
1
      type Element is private;
2
   package Data_Processing is
3
4
5
       type Data_Container (Last : Positive) is
6
         private;
7
      Data_Container_Full : exception;
8
9
      procedure Append (D : in out Data_Container;
10
                          Ε:
                                     Element);
11
12
      procedure Iterate
13
         (D : Data_Container;
14
          Proc : not null access
15
                   procedure (E : Element));
16
17
   private
18
19
20
      type Data_Container_Storage is
         array (Positive range <>) of Element;
21
22
      type Data_Container (Last : Positive) is
23
       record
24
          S
               : Data_Container_Storage (1 .. Last);
25
          Curr : Natural := 0;
26
      end record;
27
28
   end Data_Processing;
29
```

Listing 108: data_processing.adb

```
package body Data_Processing is
1
2
       procedure Append (D : in out Data_Container;
3
4
                          Ε:
                                      Element) is
       begin
5
          if D.Curr < D.S'Last then</pre>
6
             D.Curr := D.Curr + 1;
7
             D.S (D.Curr) := E;
8
          else
9
             raise Data_Container_Full;
10
             -- NOTE: This is just a dummy
11
             - -
                        implementation. A better
12
                        strategy is to add actual error
             - -
13
             - -
                        handling when the container is
14
                        full.
             - -
15
          end if;
16
       end Append;
17
18
       procedure Iterate
19
         (D : Data Container;
20
          Proc : not null access
21
                    procedure (E : Element)) is
22
```

```
23 begin
24 for I in D.S'First .. D.Curr loop
25 Proc (D.S (I));
26 end loop;
27 end Iterate;
28
29 end Data_Processing;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⇔Accessibility_Rules_Anonymous_Access_To_Subprograms.Iterator_Anonymous

MD5: fa56595ef1734f2f07ad719c36dfd8b5
```

Note that the only changes we did to the package were to remove the Process_Element type and replace the type of the Proc parameter of the Iterate procedure from a named type (Process_Element) to an anonymous type (**not null access procedure** (E : El-ement)).

Now, the same test application we used before (Show_Access_To_Subprograms) compiles as expected:

Listing 109: float data processing.ads

```
with Data_Processing;
package Float_Data_Processing is
new Data_Processing (Element => Float);
```

Listing 110: show_access_to_subprograms.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   with Float Data Processing;
3
   use Float_Data_Processing;
4
5
   procedure Show_Access_To_Subprograms is
6
7
      procedure Display (F : Float) is
8
      begin
9
          Put Line ("F :" & Float'Image (F));
10
      end Display;
11
12
      D : Data_Container (5);
13
   begin
14
        Append (D, 1.0);
15
        Append (D, 2.0);
16
        Append (D, 3.0);
17
18
        Iterate (D, Display'Access);
19
   end Show_Access_To_Subprograms;
20
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Anonymous_Access_Types.

⊶Accessibility_Rules_Anonymous_Access_To_Subprograms.Iterator_Anonymous

MD5: 64ee435aac5f2817b7d9cecf538a1e4c
```

Runtime output

- F : 1.00000E+00 F : 2.00000E+00
- F : 3.00000E+00

Remember that the compiler introduces an accessibility check in the call to Iterate, which is successful because the lifetime of Display'Access is the same as the lifetime of the Proc parameter of Iterate.

CHAPTER SEVENTEEN

LIMITED TYPES

So far, we discussed nonlimited types in most cases. In this chapter, we discuss limited types.

We can think of limited types as an easy way to avoid inappropriate semantics. For example, a lock should not be copied — neither directly, via assignment, nor with pass-by-copy. Similarly, a *file*, which is really a file descriptor, should not be copied. In this chapter, we'll see example of unwanted side-effects that arise if we don't use limited types for these cases.

17.1 Assignment and equality

Limited types have the following restrictions, which we discussed in the Introduction to Ada^{308} course:

- · copying objects of limited types via direct assignments is forbidden; and
- there's no predefined equality operator for limited types.

(Of course, in the case of nonlimited types, assignments are possible and the equality operator is available.)

By having these restrictions for limited types, we avoid inappropriate side-effects for assignment and equality operations. As an example of inappropriate side-effects, consider the case when we apply those operations on record types that have components of access types:

Listing 1: nonlimited_types.ads

```
package Nonlimited Types is
1
2
      type Simple Rec is private;
3
4
      type Integer Access is access Integer;
5
6
      function Init (I : Integer) return Simple_Rec;
7
8
      procedure Set (E : Simple_Rec;
9
                       I : Integer);
10
11
      procedure Show (E
                                : Simple Rec;
12
                        E Name : String);
13
14
   private
15
16
      type Simple_Rec is record
17
          V : Integer_Access;
18
```

(continues on next page)

³⁰⁸ https://learn.adacore.com/courses/intro-to-ada/chapters/privacy.html#intro-ada-limited-types

19 end record;
20
21 end Nonlimited_Types;

Listing 2: nonlimited_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Nonlimited_Types is
3
4
      function Init (I : Integer) return Simple_Rec
5
      is
6
      begin
7
          return E : Simple_Rec do
8
             E.V := new Integer'(I);
9
          end return;
10
      end Init;
11
12
      procedure Set (E : Simple_Rec;
13
                      I : Integer) is
14
      begin
15
         E.V.all := I;
16
      end Set;
17
18
      procedure Show (E
                            : Simple Rec;
19
                        E_Name : String) is
20
      begin
21
          Put_Line (E_Name
22
                    & ".V.all = "
23
                    & Integer'Image (E.V.all));
24
      end Show;
25
26
   end Nonlimited_Types;
27
```

Listing 3: show_wrong_assignment_equality.adb

```
1
   with Ada.Text_I0;
                            use Ada.Text_I0;
   with Nonlimited_Types; use Nonlimited_Types;
2
3
   procedure Show_Wrong_Assignment_Equality is
4
      A, B : Simple_Rec := Init (0);
5
6
7
      procedure Show_Compare is
8
      begin
         if A = B then
9
             Put_Line ("A = B");
10
          else
11
             Put_Line ("A /= B");
12
         end if;
13
      end Show_Compare;
14
   begin
15
16
      Put Line ("A := Init (0); A := Init (0);");
17
      Show (A, "A");
18
      Show (B, "B");
19
20
      Show_Compare;
      Put_Line ("-----");
21
22
      Put_Line ("Set (A, 2); Set (B, 3);");
23
      Set (A, 2);
24
```

```
Set (B, 3);
25
26
       Show (A, "A");
27
       Show (B, "B");
28
       Put_Line ("-----");
29
30
       Put_Line ("B := A");
31
       B := A;
32
33
       Show (A, "A");
Show (B, "B");
34
35
       Show_Compare;
36
       Put_Line ("-----");
37
38
       Put_Line ("Set (B, 7);");
39
       Set (B, 7);
40
41
       Show (A, "A");
42
       Show (B, "B");
43
       Show_Compare;
44
       Put Line ("-----");
45
46
   end Show_Wrong_Assignment_Equality;
47
```

Code block metadata

Runtime output

```
A := Init (0); A := Init (0);
A.V.all =
           0
B.V.all =
           0
A /= B
- - - - - - - -
Set (A, 2); Set (B, 3);
A.V.all = 2
B.V.all = 3
- - - - - - - -
B := A
A.V.all = 2
B.V.all = 2
A = B
- - - - - - - - -
Set (B, 7);
A.V.all = 7
B.V.all = 7
A = B
- - - - - - - -
```

In this code, we declare the Simple_Rec type in the Nonlimited_Types package and use it in the Show_Wrong_Assignment_Equality procedure. In principle, we're already doing many things right here. For example, we're declaring the Simple_Rec type private, so that the component V of access type is encapsulated. Programmers that declare objects of this type cannot simply mess up with the V component. Instead, they have to call the Init function and the Set procedure to initialize and change, respectively, objects of the Simple_Rec type. That being said, there are two problems with this code, which we discuss next.

The first problem we can identify is that the first call to Show_Compare shows that A and B

are different, although both have the same value in the V component (A.V.all = 0 and B.V.all = 0) — this was set by the call to the Init function. What's happening here is that the A = B expression is comparing the access values (A.V = B.V), while we might have been expecting it to compare the actual integer values after dereferencing (A.V.all = B.V.all). Therefore, the predefined equality function of the Simple_Rec type is useless and dangerous for us, as it misleads us to expect something that it doesn't do.

After the assignment of A to B (B := A), the information that the application displays seems to be correct — both A.V.all and B.V.all have the same value of two. However, when assigning the value seven to B by calling Set (B, 7), we see that the value of A.V.all has also changed. What's happening here is that the previous assignment (B := A) has actually assigned access values (B.V := A.V), while we might have been expecting it to assign the dereferenced values (B.V.all := A.V.all). Therefore, we cannot simply directly assign objects of Simple_Rec type, as this operation changes the internal structure of the type due to the presence of components of access type.

For these reasons, forbidding these operations for the Simple_Rec type is the most appropriate software design decision. If we still need assignment and equality operators, we can implement custom subprograms for the limited type. We'll discuss this topic in the next sections.

In addition to the case when we have components of access types, limited types are useful for example when we want to avoid the situation in which the same information is copied to multiple objects of the same type.

() In the Ada Reference Manual

• 7.5 Limited Types³⁰⁹

17.1.1 Assignments

Assignments are forbidden when using objects of limited types. For example:

Listing 4: limited_types.ads

```
package Limited_Types is
1
2
      type Simple Rec is limited private;
3
4
      type Integer Access is access Integer;
5
6
      function Init (I : Integer) return Simple_Rec;
7
8
   private
9
10
      type Simple Rec is limited record
11
          V : Integer_Access;
12
      end record;
13
14
   end Limited Types;
15
```

Listing 5: limited types.adb

```
1 package body Limited_Types is
2
3 function Init (I : Integer) return Simple_Rec
4 is
5 begin
```

(continues on next page)

³⁰⁹ http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

```
6 return E : Simple_Rec do
7 E.V := new Integer'(I);
8 end return;
9 end Init;
10
11 end Limited_Types;
```

Listing 6: show_limited_assignment.adb

```
with Limited_Types; use Limited_Types;
procedure Show_Limited_Assignment is
A, B : Simple_Rec := Init (0);
begin
B := A;
end Show Limited Assignment;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Assignment_

→Equality.Assignment

MD5: 019c16f7feac896fd8c37d40d0522dc8
```

Build output

In this example, we declare the limited private type Simple_Rec and two objects of this type (A and B) in the Show_Limited_Assignment procedure. (We discuss more about limited private types *later* (page 787)).

As expected, we get a compilation error for the B := A statement (in the Show_Limited_Assignment procedure). If we need to copy two objects of limited type, we have to provide a custom procedure to do that. For example, we can implement a Copy procedure for the Simple_Rec type:

Listing 7: limited types.ads

```
package Limited_Types is
1
2
      type Integer_Access is access Integer;
3
4
      type Simple Rec is limited private;
5
6
      function Init (I : Integer) return Simple_Rec;
7
8
      procedure Copy (From :
                                       Simple Rec;
9
                        To : in out Simple_Rec);
10
11
   private
12
13
      type Simple_Rec is limited record
14
         V : Integer_Access;
15
      end record;
16
17
   end Limited Types;
18
```

Listing 8: limited_types.adb

```
package body Limited_Types is
1
2
      function Init (I : Integer) return Simple_Rec
3
      is
4
      begin
5
         return E : Simple Rec do
6
             E.V := new Integer'(I);
7
         end return;
8
      end Init;
9
10
      procedure Copy (From :
                                      Simple_Rec;
11
                       To : in out Simple_Rec)
12
      is
13
      begin
14
             Copying record components
          - -
15
         To.V.all := From.V.all;
16
      end Copy;
17
18
   end Limited_Types;
19
```

Listing 9: show_limited_assignment.adb

```
with Limited_Types; use Limited_Types;
procedure Show_Limited_Assignment is
A, B : Simple_Rec := Init (0);
begin
Copy (From => A, To => B);
end Show_Limited_Assignment;
```

Code block metadata

The Copy procedure from this example copies the dereferenced values of From to To, which matches our expectation for the Simple_Rec. Note that we could have also implemented a Shallow_Copy procedure to copy the actual access values (i.e. To.V := From.V). However, having this kind of procedure can be dangerous in many case, so this design decision must be made carefully. In any case, using limited types ensures that only the assignment subprograms that are explicitly declared in the package specification are available.

17.1.2 Equality

Limited types don't have a predefined equality operator. For example:

Listing 10: limited types.ads

```
package Limited_Types is
type Integer_Access is access Integer;
type Simple_Rec is limited private;
function Init (I : Integer) return Simple_Rec;
private
(a)
```

```
10
11 type Simple_Rec is limited record
12 V : Integer_Access;
13 end record;
14
15 end Limited_Types;
```

Listing 11: limited_types.adb

```
package body Limited Types is
1
2
      function Init (I : Integer) return Simple_Rec
3
      is
4
      beain
5
          return E : Simple Rec do
6
            E.V := new Integer'(I);
7
         end return;
8
      end Init;
9
10
   end Limited_Types;
11
```

Listing 12: show limited equality.adb

```
with Ada.Text_I0;
                        use Ada.Text I0;
1
   with Limited_Types; use Limited_Types;
2
3
   procedure Show_Limited_Equality is
4
      A : Simple_Rec := Init (5);
5
      B : Simple_Rec := Init (6);
6
   begin
7
      if A = B then
8
         Put Line ("A = B");
9
      else
10
         Put_Line ("A /= B");
11
12
      end if;
  end Show_Limited_Equality;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Assignment_

→Equality.Equality

MD5: dad31b5e36de0b3b7824f723a60e5aa0
```

Build output

```
show_limited_equality.adb:8:09: error: there is no applicable operator "=" for_

_private type "Simple_Rec" defined at limited_types.ads:5
gprbuild: *** compilation phase failed
```

As expected, the comparison A = B triggers a compilation error because no predefined = operator is available for the Simple_Rec type. If we want to be able to compare objects of this type, we have to implement the = operator ourselves. For example, we can do that for the Simple_Rec type:

Listing 13: limited_types.ads

```
1 package Limited_Types is
2
3 type Integer_Access is access Integer;
4
```

```
type Simple_Rec is limited private;
5
6
       function Init (I : Integer) return Simple_Rec;
7
8
       function "=" (Left, Right : Simple_Rec)
9
                      return Boolean;
10
11
   private
12
13
       type Simple_Rec is limited record
14
         V : Integer_Access;
15
       end record;
16
17
   end Limited_Types;
18
```

Listing 14: limited_types.adb

```
package body Limited Types is
1
2
      function Init (I : Integer) return Simple_Rec
3
      is
4
      begin
5
          return E : Simple_Rec do
6
            E.V := new Integer'(I);
7
         end return;
8
      end Init;
9
10
      function "=" (Left, Right : Simple_Rec)
11
                     return Boolean is
12
      begin
13
             Comparing record components
14
         - -
         return Left.V.all = Right.V.all;
15
      end "=";
16
17
   end Limited_Types;
18
```

Listing 15: show limited equality.adb

```
use Ada.Text_I0;
  with Ada.Text_I0;
1
   with Limited_Types; use Limited_Types;
2
3
   procedure Show Limited Equality is
4
5
      A : Simple Rec := Init (5);
6
      B : Simple_Rec := Init (6);
   begin
7
      if A = B then
8
9
         Put_Line ("A = B");
10
      else
         Put_Line ("A /= B");
11
      end if;
12
  end Show_Limited_Equality;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Assignment_

→Equality.Equality

MD5: f56b2229443a5e4e33c402b41b02d318
```

Runtime output

A /= B

Here, the = operator compares the dereferenced values of Left.V and Right.V, which matches our expectation for the Simple_Rec type. Declaring types as limited ensures that we don't have unreasonable equality comparisons, and allows us to create reasonable replacements when required.

```
In other languages
In C++, you can overload the assignment operator. For example:

class Simple_Rec
{
public:
    // Overloaded assignment
    Simple_Rec& operator= (const Simple_Rec& obj);
private:
    int *V;
};
In Ada, however, we can only define the equality operator (=). Defining the assignment
operator (:=) is not possible. The following code triggers a compilation error as expected:
package Limited_Types is
    type Integer_Access is access Integer;
```

```
type Simple_Rec is limited private;
procedure ":=" (To : in out Simple Rec
```

```
From : Simple_Rec);
```

```
end Limited_Types;
```

17.2 Limited private types

As we've seen in code examples from the previous section, we can apply *information hiding* (page 38) to limited types. In other words, we can declare a type as **limited private** instead of just **limited**. For example:

```
Listing 16: simple_recs.ads
```

```
package Simple_Recs is
1
2
      type Rec is limited private;
3
4
   private
5
6
      type Rec is limited record
7
         I : Integer;
8
      end record;
9
10
  end Simple_Recs;
11
```

Code block metadata

In this case, in addition to the fact that assignments are forbidden for objects of this type (because Rec is limited), we cannot access the record components.

Note that in this example, both partial and full views of the Rec record are of limited type. In the next sections, we discuss how the partial and full views can have non-matching declarations.

In the Ada Reference Manual

```
• 7.5 Limited Types<sup>310</sup>
```

17.2.1 Non-Record Limited Types

In principle, only record types can be declared limited, so we cannot use scalar or array types. For example, the following declarations won't compile:

Listing 17: non_record_limited_error.ads

```
package Non_Record_Limited_Error is
1
2
      type Limited Enumeration is
3
        limited (Off, On);
4
5
      type Limited Integer is new
6
        limited Integer;
7
8
      type Integer_Array is
9
10
         array (Positive range <>) of Integer;
11
      type Rec is new
12
         limited Integer_Array (1 .. 2);
13
14
   end Non_Record_Limited_Error;
15
```

Code block metadata

However, we've mentioned *in a previous chapter* (page 41) that private types don't have to be record types necessarily. In this sense, limited private types makes it possible for us to use types other than record types in the full view and still benefit from the restrictions of limited types. For example:

Listing 18: simple recs.ads

```
1 package Simple_Recs is
2
3 type Limited_Enumeration is
4 limited private;
5
6 type Limited_Integer is
```

(continues on next page)

³¹⁰ http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

```
limited private;
7
8
       type Limited Integer Array 2 is
9
         limited private;
10
11
   private
12
13
       type Limited Enumeration is (Off, On);
14
15
       type Limited_Integer is new Integer;
16
17
       type Integer_Array is
18
         array (Positive range <>) of Integer;
19
20
       type Limited_Integer_Array_2 is
21
         new Integer_Array (1 .. 2);
22
23
   end Simple_Recs;
24
```

Code block metadata

Here, Limited_Enumeration, Limited_Integer, and Limited_Integer_Array_2 are limited private types that encapsulate an enumeration type, an integer type, and a constrained array type, respectively.

17.2.2 Partial and full view of limited types

In the previous example, both partial and full views of the Rec type were limited. We may actually declare a type as **limited private** (in the public part of a package), while its full view is nonlimited. For example:

Listing 19: simple_recs.ads

```
package Simple Recs is
1
2
      type Rec is limited private;
3
       -- Partial view of Rec is limited
4
5
   private
6
7
      type Rec is record
8
       -- Full view of Rec is nonlimited
9
         I : Integer:
10
      end record;
11
12
   end Simple Recs;
13
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Private_
GTypes.Limited_Partial_Full_View
MD5: 5d0dbc3e87531476856f0ac1f9b22c78
```

In this case, only the partial view of Rec is limited, while its full view is nonlimited. When deriving from Rec, the view of the derived type is the same as for the parent type:

Listing 20: simple_recs-child.ads

```
package Simple Recs.Child
1
   is
2
      type Rec Derived is new Rec;
3
      -- As for its parent, the
4
      -- partial view of Rec_Derived
5
      -- is limited, but the full view
-- is nonlimited.
6
7
8
  end Simple Recs.Child;
9
```

Code block metadata

Clients must nevertheless comply with their partial view, and treat the type as if it is in fact limited. In other words, if you use the Rec type in a subprogram or package outside of the Simple_Recs package (or its child packages), the type is limited from that perspective:

Listing 21: use_rec_in_subprogram.adb

```
with Simple_Recs; use Simple_Recs;
procedure Use_Rec_In_Subprogram is
R1, R2 : Rec;
begin
R1.I := 1;
R2 := R1;
end Use_Rec_In_Subprogram;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Private_

GTypes.Limited_Partial_Full_View

MD5: f0af323a951853b97a2b67ce9b13e732
```

Build output

Here, compilation fails because the type Rec is limited from the procedure's perspective.

Limitations

Note that the opposite — declaring a type as **private** and its full full view as **limited private** — is not possible. For example:

Listing 22: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is private;
4
5 private
6
```

```
7 type Rec is limited record
8 I : Integer;
9 end record;
10
11 end Simple_Recs;
```

Code block metadata

Build output

```
simple_recs.ads:7:09: error: completion of nonlimited type cannot be limited
gprbuild: *** compilation phase failed
```

As expected, we get a compilation error in this case. The issue is that the partial view cannot be allowed to mislead the client about what's possible. In this case, if the partial view allows assignment, then the full view must actually provide assignment. But the partial view can restrict what is actually possible, so a limited partial view need not be completed in the full view as a limited type.

In addition, tagged limited private types cannot have a nonlimited full view. For example:

Listing 23: simple_recs.ads

```
package Simple Recs is
1
2
       type Rec is tagged limited private;
3
4
   private
5
6
       type Rec is tagged record
7
          I : Integer;
8
      end record:
9
10
   end Simple Recs;
11
```

Code block metadata

Build output

Here, compilation fails because the type Rec is nonlimited in its full view.

17.2.3 Limited and nonlimited in full view

Declaring the full view of a type as limited or nonlimited has implications in the way we can use objects of this type in the package body. For example: Listing 24: simple_recs.ads

```
package Simple Recs is
1
2
      type Rec_Limited_Full is limited private;
3
      type Rec_Nonlimited_Full is limited private;
4
5
      procedure Copy
6
         (From :
                         Rec Limited Full;
7
         To : in out Rec Limited Full);
8
      procedure Copy
9
                         Rec_Nonlimited_Full;
         (From :
10
         To : in out Rec_Nonlimited_Full);
11
12
   private
13
14
       type Rec_Limited_Full is limited record
15
         I : Integer;
16
      end record;
17
18
       type Rec Nonlimited Full is record
19
         I : Integer;
20
       end record;
21
22
   end Simple_Recs;
23
```

Listing 25:	simple	recs.adb
-------------	--------	----------

```
package body Simple_Recs is
1
2
      procedure Copy
3
         (From :
                        Rec_Limited_Full;
4
          To : in out Rec_Limited_Full)
5
      is
6
      begin
7
         To := From;
8
          -- ERROR: assignment is forbidden because
9
         - -
                    Rec_Limited_Full is limited in
10
                     its full view.
         - -
11
      end Copy;
12
13
      procedure Copy
14
        (From :
                        Rec_Nonlimited_Full;
15
          To : in out Rec_Nonlimited_Full)
16
      is
17
      begin
18
         To := From;
19
          -- OK: assignment is allowed because
20
                  Rec_Nonlimited_Full is
          - -
21
                  nonlimited in its full view.
          - -
22
      end Copy;
23
24
   end Simple_Recs;
25
```

Code block metadata

Build output

```
simple_recs.adb:8:07: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed
```

Here, both Rec_Limited_Full and Rec_Nonlimited_Full are declared as **private limited**. However, Rec_Limited_Full type is limited in its full view, while Rec_Nonlimited_Full is nonlimited. As expected, the compiler complains about the To := From assignment in the Copy procedure for the Rec_Limited_Full type because its full view is limited (so no assignment is possible). Of course, in the case of the objects of Rec_Nonlimited_Full type, this assignment is perfectly fine.

17.2.4 Limited private component

Another example mentioned by the Ada Reference Manual $(7.3.1^{311}, 5/1)$ is about an array type whose component type is limited private, but nonlimited in its full view. Let's see a complete code example for that:

```
Listing 26: limited nonlimited arrays.ads
```

```
package Limited_Nonlimited_Arrays is
1
2
      type Limited Private is
3
         limited private;
4
5
      function Init return Limited_Private;
6
7
      -- The array type Limited_Private_Array
8
      -- is limited because the type of its
9
       -- component is limited.
10
      type Limited_Private_Array is
11
         array (Positive range <>) of
12
           Limited Private;
13
14
   private
15
16
      type Limited_Private is
17
       record
18
         A : Integer;
19
      end record;
20
21
       -- Limited_Private_Array type is
22
      - -
          nonlimited at this point because
23
       - -
          its component is nonlimited.
24
25
          The assignments below are OK:
26
      A1 : Limited_Private_Array (1 .. 5);
27
28
      A2 : Limited_Private_Array := A1;
29
30
   end Limited_Nonlimited_Arrays;
31
```

Listing 27: limited nonlimited arrays.adb

```
package body Limited_Nonlimited_Arrays is
function Init return Limited_Private is
((A => 1));
end Limited_Nonlimited_Arrays;
```

³¹¹ http://www.ada-auth.org/standards/22rm/html/RM-7-3-1.html

Listing 28: show_limited_nonlimited_array.adb

```
with Limited Nonlimited Arrays;
1
   use Limited_Nonlimited_Arrays;
2
3
   procedure Show Limited Nonlimited Array is
4
      A3 : Limited Private Array (1 .. 2) :=
5
              (others => Init);
6
      A4 : Limited_Private_Array (1 .. 2);
7
   begin
8
      -- ERROR: this assignment is illegal because
9
      -- Limited_Private_Array is limited, as
10
      -- its component is limited at this point.
11
      A4 := A3:
12
  end Show Limited Nonlimited Array;
13
```

Code block metadata

Build output

As we can see in this example, the limitedness of the array type Limited_Private_Array depends on the limitedness of its component type Limited_Private. In the private part of Limited_Nonlimited_Arrays package, where Limited_Private is nonlimited, the array type Limited_Private_Array becomes nonlimited as well. In contrast, in the Show_Limited_Nonlimited_Array, the array type is limited because its component is limited in that scope.

In the Ada Reference Manual

• 7.3.1 Private Operations³¹²

17.2.5 Tagged limited private types

For tagged private types, the partial and full views must match: if a tagged type is limited in the partial view, it must be limited in the full view. For example:

Listing 29: simple recs.ads

```
package Simple_Recs is
type Rec is tagged limited private;
type Rec is tagged limited record
type Rec is tagged limited record
I : Integer;
end record;
```

(continues on next page)

³¹² http://www.ada-auth.org/standards/22rm/html/RM-7-3-1.html

11 end Simple_Recs;

10

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Private_

→Types.Tagged_Limited_Private_Types

MD5: bee48bd7e0d70ddfd288c0de5e21b039
```

Here, the tagged Rec type is limited both in its partial and full views. Any mismatch in one of the views triggers a compilation error. (As an exercise, you may remove any of the **limited** keywords from the code example and try to compile it.)

For further reading....

This rule is for the sake of dynamic dispatching and classwide types. The compiler must not allow any of the types in a derivation class — the set of types related by inheritance to be different regarding assignment and equality (and thus inequality). That's necessary because we are meant to be able to manipulate objects of any type in the entire set of types via the partial view presented by the root type, without knowing which specific tagged type is involved.

17.3 Explicitly limited types

Under certain conditions, limited types can be called explicitly limited — note that using the **limited** keyword in a part of the declaration doesn't necessary ensure this, as we'll see later.

Let's start with an example of an explicitly limited type:

Listing 30: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is limited record
4 I : Integer;
5 end record;
6
7 end Simple_Recs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Explicitly_Limited_
GTypes.Explicitly_Limited_Types
MD5: de73a20140628420830ed9fe0b2dedb5
```

The Rec type is also explicitly limited when it's declared limited in the private type's completion (in the package's private part):

Listing 31: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is limited private;
4
5 private
6
```

```
7 type Rec is limited record
8 I : Integer;
9 end record;
10
11 end Simple_Recs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Explicitly_Limited_

GTypes.Explicitly_Limited_Types

MD5: ececb364f5365a74db43952e9421dee0
```

In this case, Rec is limited both in the partial and in the full view, so it's considered explicitly limited.

However, as we've learned before (page 789), we may actually declare a type as **limited private** in the public part of a package, while its full view is nonlimited. In this case, the limited type is not considered explicitly limited anymore.

For example, if we make the full view of the Rec nonlimited (by removing the **limited** keyword in the private part), then the Rec type isn't explicitly limited anymore:

Listing 32: simple_recs.ads

```
package Simple_Recs is
1
2
       type Rec is limited private;
3
4
   private
5
6
       type Rec is record
7
          I : Integer;
8
       end record;
9
10
   end Simple_Recs;
11
```

Code block metadata

Now, even though the Rec type was declared as limited private, the full view indicates that it's actually a nonlimited type, so it isn't explicitly limited.

Note that *tagged limited private types* (page 794) are always explicitly limited types — because, as we've learned before, they cannot have a nonlimited type declaration in its full view.

In the Ada Reference Manual

- 6.2 Formal Parameter Modes³¹³
- 6.4.1 Parameter Associations³¹⁴
- 7.5 Limited Types³¹⁵

³¹³ http://www.ada-auth.org/standards/22rm/html/RM-6-2.html

³¹⁴ http://www.ada-auth.org/standards/22rm/html/RM-6-4-1.html

³¹⁵ http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

17.4 Subtypes of Limited Types

We can declare subtypes of limited types. For example:

```
Listing 33: simple recs.ads
   package Simple Recs is
1
2
      type Limited Integer Array (L : Positive) is
3
         limited private;
4
5
      subtype Limited Integer Array 2 is
6
         Limited_Integer_Array (2);
8
   private
9
10
      type Integer Array is
11
         array (Positive range <>) of Integer;
12
13
      type Limited Integer Array (L : Positive) is
14
         limited record
15
         Arr : Integer Array (1 .. L);
16
      end record;
17
18
   end Simple Recs;
19
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_

→Limited_Types.Limited_Subtype

MD5: 2a82c3c96fad2a01b9a8c15912d4b974
```

Here, Limited_Integer_Array_2 is a subtype of the Limited_Integer_Array type. Since Limited_Integer_Array is a limited type, the Limited_Integer_Array_2 subtype is limited as well. A subtype just introduces a name for some constraints on an existing type. As such, a subtype doesn't change the limitedness of the constrained type.

We can test this in a small application:

Listing 34: test_limitedness.adb

```
with Simple_Recs; use Simple_Recs;
procedure Test_Limitedness is
Dummy_1, Dummy_2 : Limited_Integer_Array_2;
begin
Dummy_2 := Dummy_1;
end Test_Limitedness;
```

Code block metadata

Build output

```
test_limitedness.adb:6:04: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed
```

As expected, compilations fails because Limited_Integer_Array_2 is a limited (sub)type.

17.5 Deriving from limited types

In this section, we discuss the implications of deriving from limited types. As usual, let's start with a simple example:

Listing 35: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is limited null record;
4
5 type Rec_Derived is new Rec;
6
7 end Simple_Recs;
```

Code block metadata

In this example, the Rec_Derived type is derived from the Rec type. Note that the Rec_Derived type is limited because its ancestor is limited, even though the **limited** keyword doesn't show up in the declaration of the Rec_Derived type. Note that we could have actually used the **limited** keyword here:

type Rec_Derived is limited new Rec;

Therefore, we cannot use the assignment operator for objects of Rec_Derived type:

Listing 36: test_limitedness.adb

```
with Simple_Recs; use Simple_Recs;
procedure Test_Limitedness is
Dummy_1, Dummy_2 : Rec_Derived;
begin
Dummy_2 := Dummy_1;
end Test_Limitedness;
```

Code block metadata

Build output

```
test_limitedness.adb:6:04: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed
```

Note that we cannot derive a limited type from a nonlimited ancestor:

Listing 37: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is null record;
4
5 type Rec_Derived is limited new Rec;
6
7 end Simple_Recs;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_

→Limited_Types.Derived_Limited_Type_Nonlimited_Ancestor

MD5: 78a7574cc6233ddc826359acb6e644ee
```

Build output

```
simple_recs.ads:5:04: error: parent type "Rec" of limited type must be limited
gprbuild: *** compilation phase failed
```

As expected, the compiler indicates that the ancestor Rec should be of limited type.

In fact, all types in a derivation class are the same — either limited or not. (That is especially important with dynamic dispatching via tagged types. We discuss this topic in another chapter.)

1 In the Ada Reference Manual

- 7.3 Private Types and Private Extensions³¹⁶
- 7.5 Limited Types³¹⁷

17.5.1 Deriving from limited private types

Of course, we can also derive from limited private types. However, there are more rules in this case than the ones we've seen so far. Let's start with an example:

```
Listing 38: simple_recs.ads
```

```
1 package Simple_Recs is
2
3 type Rec is limited private;
4
5 private
6
7 type Rec is limited null record;
8
9 end Simple_Recs;
```

Listing 39: simple recs-ext.ads

```
package Simple Recs.Ext is
1
2
       type Rec Derived is new Rec;
3
4
          OR:
       - -
5
       - -
6
           type Rec Derived is
       - -
7
              limited new Rec;
       - -
8
9
   end Simple Recs.Ext;
10
```

Listing 40: test limitedness.adb

```
with Simple_Recs.Ext; use Simple_Recs.Ext;
```

2

³¹⁶ http://www.ada-auth.org/standards/22rm/html/RM-7-3.html ³¹⁷ http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

```
3 procedure Test_Limitedness is
4 Dummy_1, Dummy_2 : Rec_Derived;
5 begin
6 Dummy_2 := Dummy_1;
7 end Test_Limitedness;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_

→Limited_Types.Derived_Limited_Private_Type

MD5: c6eed14520589b9c1e11c17bd6179c19
```

Build output

```
test_limitedness.adb:6:04: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed
```

Here, Rec_Derived is a limited type derived from the (limited private) Rec type. We can verify that Rec_Derived type is limited because the compilation of the Test_Limitedness procedure fails.

17.5.2 Deriving from non-explicitly limited private types

Up to this point, we have discussed *explicitly limited types* (page 795). Now, let's see how derivation works with *non-explicitly* limited types.

Any type derived from a limited type is always limited, even if the full view of its ancestor is nonlimited. For example, let's modify the full view of Rec and make it nonlimited (i.e. make it *not explicitly* limited):

Listing 41: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is limited private;
4
5 private
6
7 type Rec is null record;
8
9 end Simple_Recs;
```

Code block metadata

Build output

```
simple_recs.ads:1: Simple_Recs cannot be used as a main program
gprbind: invocation of gnatbind failed
gprbuild: unable to bind simple_recs.ads
```

Here, Rec_Derived is a limited type because the partial view of Rec is limited. The fact that the full view of Rec is nonlimited doesn't affect the Rec_Derived type — as we can verify with the compilation error in the Test_Limitedness procedure.

Note, however, that a derived type becomes nonlimited in the **private part or the body** of a child package if it isn't explicitly limited. In this sense, the derived type inherits the *nonlimitedness* of the parent's full view. For example, because we're declaring Rec_Derived as **is**

new Rec in the child package (Simple_Recs.Ext), we're saying that Rec_Derived is limited *outside* this package, but nonlimited in the private part and body of the Simple_Recs.Ext package. We can verify this by copying the code from the Test_Limitedness procedure to a new procedure in the body of the Simple_Recs.Ext package:

Listing 42: simple_recs-ext.ads

```
1
   package Simple Recs.Ext
2
     with Elaborate_Body is
3
     -- Rec_Derived is derived from Rec, which is a
4
     -- limited private type that is nonlimited in
5
     -- its full view.
6
7
     -- Rec Derived isn't explicitly limited.
8
     -- Therefore, it's nonlimited in the private
9
     -- part of Simple_Recs.Ext and its package
10
11
     -- body.
     - -
12
13
     type Rec_Derived is new Rec;
14
  end Simple_Recs.Ext;
15
```

Listing 43: simple_recs-ext.adb

```
package body Simple_Recs.Ext is
1
2
      procedure Test Child Limitedness is
3
         Dummy_1, Dummy_2 : Rec_Derived;
4
      begin
5
             Here, Rec Derived is a nonlimited
          - -
6
          -- type because Rec is nonlimited in
7
          -- its full view.
8
9
         Dummy 2 := Dummy 1;
10
      end Test_Child_Limitedness;
11
12
   end Simple Recs.Ext;
13
```

Listing 44: test_limitedness.adb

```
-- We copied the code to the
1
       Test_Child_Limitedness procedure (in the
   - -
2
       body of the Simple_Recs.Ext package) and
   - -
3
   - -
       commented it out here.
4
   - -
5
       You may uncomment the code to verify
   - -
6
   - -
       that Rec Derived is limited in this
7
   - -
       procedure.
8
   - -
9
10
   -- with Simple_Recs.Ext; use Simple_Recs.Ext;
11
12
   procedure Test_Limitedness is
13
      -- Dummy_1, Dummy_2 : Rec_Derived;
14
   begin
15
          Dummy_2 := Dummy_1;
      - -
16
      null;
17
   end Test_Limitedness;
18
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_ →Limited_Types.Derived_Limited_Private_Type MD5: f480cd05afff622e451684a0293cb982

In the Test_Child_Limitedness procedure of the Simple_Recs.Ext package, we can use the Rec_Derived as a nonlimited type because its ancestor Rec is nonlimited in its full view. (*As we've learned before* (page 791), if a limited type is nonlimited in its full view, we can copy objects of this type in the private part of the package specification or in the package body.)

Outside of the package, both Rec and Rec_Derived types are limited types. Therefore, if we uncomment the code in the Test_Limitedness procedure, compilation fails there (because Rec_Derived is viewed as descending from a limited type).

Deriving from tagged limited private types

The rules for deriving from tagged limited private types are slightly different than the rules we've seen so far. This is because tagged limited types are always *explicitly limited types* (page 795).

Let's look at an example:

```
Listing 45: simple recs.ads
```

```
package Simple_Recs is
type Tagged_Rec is tagged limited private;
private
type Tagged_Rec is tagged limited null record;
end Simple_Recs;
```

Listing 46: simple_recs-ext.ads

```
package Simple Recs.Ext is
1
2
      type Rec Derived is new
3
         Tagged Rec with private;
4
5
   private
6
7
      type Rec Derived is new
8
        Tagged Rec with null record;
9
10
   end Simple Recs.Ext;
11
```

Listing 47: test limitedness.adb

```
with Simple_Recs.Ext; use Simple_Recs.Ext;
procedure Test_Limitedness is
Dummy_1, Dummy_2 : Rec_Derived;
begin
Dummy_2 := Dummy_1;
end Test_Limitedness;
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_ →Limited_Types.Derived_Tagged_Limited_Private_Type MD5: 81c8a010f093d8823b84bb6e69c4114e

Build output

test_limitedness.adb:6:04: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed

In this example, Rec_Derived is a tagged limited type derived from the Tagged_Rec type. (Again, we can verify the limitedness of the Rec_Derived type with the Test_Limitedness procedure.)

As explained previously, the derived type (Rec_Derived) is a limited type, even though the **limited** keyword doesn't appear in its declaration. We could, of course, include the **limited** keyword in the declaration of Rec Derived:

Listing 48: simple_recs-ext.ads

```
package Simple Recs.Ext is
1
2
      type Rec Derived is limited new
3
         Tagged Rec with private;
4
5
   private
6
7
      type Rec Derived is limited new
8
         Tagged Rec with null record;
9
10
   end Simple Recs.Ext;
11
```

Code block metadata

Build output

simple_recs-ext.ads:1: Simple_Recs.ext cannot be used as a main program
gprbind: invocation of gnatbind failed
gprbuild: unable to bind simple_recs-ext.ads

(Obviously, if we include the **limited** keyword in the partial view of the derived type, we must include it in its full view as well.)

Deriving from limited interfaces

The rules for limited interfaces are different from the ones for limited tagged types. In contrast to the rule we've seen in the previous section, a type that is derived from a limited type isn't automatically limited. In other words, it does **not** inherit the *limitedness* from the interface. For example:

Listing 49: simple recs.ads

```
1 package Simple_Recs is
2
3 type Limited_IF is limited interface;
4
5 end Simple Recs;
```

```
Listing 50: simple_recs-ext.ads
```

```
package Simple_Recs.Ext is
1
2
      type Rec Derived is new
3
         Limited_IF with private;
4
5
   private
6
7
      type Rec_Derived is new
8
         Limited_IF with null record;
9
10
   end Simple_Recs.Ext;
11
```

Listing 51: test_limitedness.adb

```
with Simple_Recs.Ext; use Simple_Recs.Ext;
procedure Test_Limitedness is
Dummy_1, Dummy_2 : Rec_Derived;
begin
Dummy_2 := Dummy_1;
end Test_Limitedness;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_

→Limited_Types.Derived_Interface_Limited_Private

MD5: d9cf0bd26b86d0caec82eff2a2ec6ead
```

Here, Rec_Derived is derived from the limited Limited_IF interface. As we can see, the Test_Limitedness compiles fine because Rec_Derived is nonlimited.

Of course, if we want Rec_Derived to be limited, we can make this explicit in the type declaration:

Listing 52: simple_recs-ext.ads

```
package Simple_Recs.Ext is
1
2
      type Rec_Derived is limited new
3
         Limited_IF with private;
4
5
   private
6
7
      type Rec_Derived is limited new
8
         Limited_IF with null record;
9
10
   end Simple_Recs.Ext;
11
```

Listing 53: test_limitedness.adb

```
with Simple_Recs.Ext; use Simple_Recs.Ext;
procedure Test_Limitedness is
Dummy_1, Dummy_2 : Rec_Derived;
begin
Dummy_2 := Dummy_1;
rend Test_Limitedness;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Deriving_From_

→Limited_Types.Derived_Interface_Limited_Private

MD5: abb295cbfd5ade5f351991c2fbaf519c
```

Build output

```
test_limitedness.adb:6:04: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed
```

Now, compilation of Test_Limitedness fails because Rec_Derived is explicitly limited.

17.6 Immutably Limited Types

According to the Annotated Ada Reference Manual (7.5, 8.b/3)³¹⁸, "an immutably limited type is a type that cannot become nonlimited subsequently in a private part or in a child unit." In fact, while we were talking about *partial and full view of limited types* (page 789), we've seen that limited private types can become nonlimited in their full view. Such limited types are *not* immutably limited.

The Annotated Ada Reference Manual also says that "if a view of the type makes it immutably limited, then no copying (assignment) operations are ever available for objects of the type. This allows other properties; for instance, it is safe for such objects to have access discriminants that have defaults or designate other limited objects." We'll see examples of this later on.

Immutably limited types include:

- explicitly limited types (page 795)
- tagged limited types (i.e. with the keywords tagged limited);
- tagged limited private types (page 794);
- limited private type that have at least one access discriminant (page 725) with a default expression;
- task types, protected types, and synchronized interfaces;
- any types derived from immutably limited types.

Let's look at a code example that shows instances of immutably limited types:

Listing 54: show_immutably_limited_types.ads

```
package Show Immutably Limited Types is
1
2
3
       -- Explicitly limited type
4
5
       type Explicitly_Limited_Rec is limited
6
       record
7
         A : Integer;
8
      end record;
9
10
11
       - -
          Tagged limited type
12
13
       type Limited Tagged Rec is tagged limited
14
       record
15
          A : Integer;
16
      end record;
17
```

(continues on next page)

³¹⁸ http://www.ada-auth.org/standards/22aarm/html/AA-7-5.html

18

(continued from previous page)

```
19
       - -
          Tagged limited private type
20
21
       type Limited_Tagged_Private is
22
         tagged limited private;
23
24
25
       - -
       -- Limited private type with an access
26
       - -
          discriminant that has a default
27
       - -
          expression
28
29
       type Limited_Rec_Access_D
30
         (AI : access Integer := new Integer) is
31
           limited private;
32
33
34
       - -
       -- Task type
35
36
       task type TT is
37
        entry Start;
38
        entry Stop;
39
       end TT;
40
41
42
       - -
       -- Protected type
43
44
       - -
       protected type PT is
45
        function Value return Integer;
46
       private
47
       A : Integer;
48
       end PT;
49
50
51
      -- Synchronized interface
52
53
      type SI is synchronized interface;
54
55
56
      -- A type derived from an immutably
57
      -- limited type
58
59
      type Derived_Immutable is new
60
        Explicitly_Limited_Rec;
61
62
   private
63
64
       type Limited_Tagged_Private is tagged limited
65
       record
66
         A : Integer;
67
       end record;
68
69
       type Limited Rec Access D
70
         (AI : access Integer := new Integer)
71
       is limited
72
        record
73
           A : Integer;
74
         end record;
75
76
   end Show_Immutably_Limited_Types;
77
```

```
Listing 55: show_immutably_limited_types.adb
```

```
package body Show_Immutably_Limited_Types is
1
2
      task body TT is
3
      begin
4
         accept Start;
5
         accept Stop;
6
      end TT;
7
8
      protected body PT is
9
         function Value return Integer is
10
           (PT.A);
11
      end PT:
12
13
   end Show Immutably Limited Types;
14
```

Code block metadata

Build output

```
show_immutably_limited_types.ads:31:30: warning: coextension will not be_
_deallocated when its associated owner is deallocated [enabled by default]
```

In the Show_Immutably_Limited_Types package above, we see multiple instances of immutably limited types. (The comments in the source code indicate each type.)

1 In the Ada Reference Manual

• 7.5 Limited Types³¹⁹

17.6.1 Non immutably limited types

Not every limited type is immutably limited. We already mentioned untagged private limited types, which can *become nonlimited in their full view* (page 789). In addition, we have nonsynchronized limited interface types. As mentioned earlier in this chapter, a *type derived from a nonsynchronized limited interface* (page 803), can be nonlimited, so it's not immutably limited.

1 In the Ada Reference Manual

- 7.3.1 Private Operations³²⁰
- 7.5 Limited Types³²¹

³¹⁹ http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

³²⁰ http://www.ada-auth.org/standards/22rm/html/RM-7-3-1.html

³²¹ http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

17.7 Limited Types with Discriminants

In this section, we look into the implications of using discriminants with limited types. Actually, most of the topics mentioned here have already been covered in different sections of previous chapters, as well as in this chapter. Therefore, this section is in most parts just a review of what we've already discussed.

Let's start with a simple example:

Listing 56: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec (L : Positive)
4 is limited null record;
5
6 end Simple_Recs;
```

Listing 57: test_limitedness.adb

```
with Simple Recs; use Simple Recs;
1
2
   procedure Test Limitedness is
3
      Dummy_1 : Rec (2);
4
      Dummy_2 : Rec (3);
5
   begin
6
      Dummy_2 := Dummy_1;
7
              ~~~~~
8
       - -
       -- ERRORS:
9
           1. Cannot assign objects of
       - -
10
                limited types.
       - -
11
       - -
            2. Cannot assign objects with
12
               different discriminants.
13
  end Test_Limitedness;
14
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Discriminants.

→Simple_Example

MD5: 7b4a62c0341becf16f59e163b4359397
```

Build output

```
test_limitedness.adb:7:04: error: left hand of assignment must not be limited type
gprbuild: *** compilation phase failed
```

In this example, we see the declaration of the limited type Rec, which has the discriminant L. For objects of type Rec, we not only have the typical restrictions that *equality and assignment aren't available* (page 782), but we also have the restriction that we won't be able to assign objects with different discriminants.

1 In the Ada Reference Manual

3.7 Discriminants³²²

```
322 http://www.ada-auth.org/standards/12rm/html/RM-3-7.html
```

17.7.1 Default Expressions

On the other hand, there are restrictions that apply to nonlimited types with discriminants, but not to limited types with discriminants. This concerns mostly default expressions, which are generally allowed for discriminants of limited types.

Discriminants of tagged limited types

As we've discussed previously, we can use default expressions for discriminants of tagged limited types. Let's see an example:

Listing 58: recs.ads

```
1 package Recs is
2
3 type LTT (L : Positive := 1;
4 M : Positive := 2) is
5 tagged limited null record;
6
7 end Recs;
```

Code block metadata

Obviously, the same applies to tagged limited private types (page 794):

```
Listing 59: recs.ads
```

```
package Recs is
1
2
       type LTT (L : Positive := 1;
3
                 M : Positive := 2) is
4
         tagged limited private;
5
6
   private
7
8
       type LTT (L : Positive := 1;
9
                 M : Positive := 2) is
10
         tagged limited null record;
11
12
   end Recs;
13
```

Code block metadata

In the case of tagged, nonlimited types, using default expressions in this context isn't allowed.

Access discriminant

Similarly, when using limited types, we can specify default expressions for *access discriminants* (page 725):

Listing 60: custom_recs.ads

```
package Custom Recs is
1
2
      -- Specifying a default expression for
3
      -- an access discriminant:
4
      type Rec (IA : access Integer :=
5
                        new Integer'(0)) is limited
6
      record
7
         I : Integer := IA.all;
8
      end record;
9
10
   end Custom Recs;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Discriminants.

⇔Access_Discriminant_Default_Expression

MD5: 23703d9dc80e9f1c8fe237c76b9dd6b0
```

Build output

```
custom_recs.ads:6:21: warning: coextension will not be deallocated when its...

associated owner is deallocated [enabled by default]
```

In fact, as we've discussed before (page 727), this isn't possible for nonlimited types.

Note, however, that we can only assign a default expression to an access discriminant of an *immutably limited type* (page 805).

Discriminants of nontagged limited types

In addition to tagged limited types, we can use default expressions for discriminants of nontagged limited types. Let's see an example:

Listing 61: recs.ads

```
1 package Recs is
2
3 type LTT (L : Positive := 1;
4 M : Positive := 2) is
5 limited null record;
6
7 end Recs;
```

Code block metadata

Obviously, the same applies to *limited private types* (page 787):

Listing 62: recs.ads

```
1 package Recs is
2
3 type LTT (L : Positive := 1;
4 M : Positive := 2) is
5 limited private;
6
```

```
7 private
8
9 type LTT (L : Positive := 1;
10 M : Positive := 2) is
11 limited null record;
12
13 end Recs;
```

Code block metadata

Note that using default expressions for discriminants of nonlimited, nontagged types is OK as well.

Mutable subtypes and Limitedness

As we've mentioned before, an unconstrained discriminated subtype with defaults is called a mutable subtype. An important feature of mutable subtypes is that it allows changing the discriminants of an object, e.g. via assignments. However, as we know, we cannot assign to objects of limited types. Therefore, in essence, a type should be nonlimited to be considered a mutable subtype.

Let's look at a code example:

```
Listing 63: recs.ads
```

```
package Recs is
1
2
       type LTT (L : Positive := 1;
3
                 M : Positive := 2) is
4
         limited null record;
5
6
      function Init (L : Positive;
7
                       M : Positive
8
                       return LTT is
9
         ((L => L, M => M));
10
11
       procedure Copy (From :
                                       LTT;
12
13
                        To : in out LTT);
14
15
   end Recs;
```

```
Listing 64: recs.adb
```

```
package body Recs is
1
2
      procedure Copy (From :
                                      LTT;
3
                        To : in out LTT) is
4
      beain
5
          To := Init (L => From.L,
6
                       M \implies From.M);
7
              ERROR: cannot assign to object of
8
                      limited type
          - -
9
10
          To.L := From.L;
11
          To.M := From.M;
12
          -- ERROR: cannot change discriminants
13
```

14 end Copy;
15
16 end Recs;

Listing 65: show.adb

```
1 with Recs; use Recs;
2
3 procedure Show is
4 A : LTT;
5 B : LTT := Init (10, 12);
6 begin
7 Copy (From => B, To => A);
8 end Show;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Discriminants.

→Discriminant_Default_Value_Tagged_TYpe

MD5: e8dfble99e33923aa4023428ecb17372
```

Build output

```
recs.adb:6:07: error: left hand of assignment must not be limited type
recs.adb:11:09: error: assignment to discriminant not allowed
recs.adb:12:09: error: assignment to discriminant not allowed
gprbuild: *** compilation phase failed
```

As we can see in the Copy procedure, it's not possible to properly assign to the target object. Using Init is forbidden because the assignment is not initializing the target object — as we're not declaring To at this point. Also, changing the individual discriminants is forbidden as well. Therefore, we don't have any means to change the discriminants of the target object. (In contrast, if LTT was a nonlimited type, we would be able to implement Copy by using the call to the Init function.)

17.7.2 Limited private type with unknown discriminants

We can declare limited private types with *unknown discriminants* (page 217). Let's see an example:

Listing 66: limited_private_unknown_discriminants.ads

```
package Limited Private Unknown Discriminants is
1
2
      type Rec (<>) is limited private;
3
4
   private
5
6
      type Rec is limited
7
      record
8
         I : Integer;
9
      end record;
10
11
  end Limited Private Unknown Discriminants;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Discriminants.

→Limited_Private_Unknown_Discriminants

MD5: 74184919132a084da76bd3e1445c22e5
```

(continued from previous page)

In this example, we declare type Rec, which has unknown discriminants.

As we mentioned earlier, when we use a private type with unknown discriminants, we gain extra control over its initialization. In addition, if we declare those types as limited, we gain even more control. In fact, this is what the Annotated Ada Reference Manual (3.7, 26.b/2)³²³ says:

"A subtype with unknown discriminants is indefinite, and hence an object of such a subtype needs explicit initialization. A limited private type with unknown discriminants is 'extremely' limited; objects of such a type can be initialized only by subprograms (either procedures with a parameter of the type, or a function returning the type) declared in the package. Subprograms declared elsewhere can operate on and even return the type, but they can only initialize the object by calling (ultimately) a subprogram in the package declaring the type. Such a type is useful for keeping complete control over object creation within the package declaring the type."

Let's reuse a code example from the *previous section on unknown discriminants* (page 219) and use limited types:

Listing 67: limited private unknown discriminants.ads

```
package Limited Private Unknown Discriminants is
1
2
       type Rec (<>) is limited private;
3
4
       function Init return Rec;
5
6
   private
7
8
       type Rec is limited
9
       record
10
          I : Integer;
11
       end record;
12
13
       function Init return Rec is
14
         ((I => 0));
15
16
   end Limited Private Unknown Discriminants;
17
```

Listing 68: show_constructor_function.adb

```
with Limited_Private_Unknown_Discriminants;
use Limited_Private_Unknown_Discriminants;
procedure Show_Constructor_Function is
        R : Rec := Init;
        begin
        null;
        end Show_Constructor_Function;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Discriminants.

→Limited_Private_Unknown_Discriminants

MD5: f4b1de2a83837e2e52b0b57214f0eaf9
```

A function such as Init is called a *constructor function for limited types* (page 819). We discuss this topic in more detail later on.

³²³ http://www.ada-auth.org/standards/22aarm/html/AA-3-7.html

17.8 Record components of limited type

In this section, we discuss the implications of using components of limited type. Let's start by declaring a record component of limited type:

Listing 69: simple_recs.ads

```
package Simple_Recs is
1
2
       type Int Rec is limited record
3
         V : Integer:
4
      end record;
5
6
       type Rec is limited record
7
          IR : Int Rec;
8
      end record;
9
10
   end Simple_Recs;
11
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Record_Components_
Limited_Type.Record_Components_Limited_Type
MD5: 71badd1e38cc4ff37f16d99dd203614b
```

As soon as we declare a record component of some limited type, the whole record is limited. In this example, the Rec record is limited due to the presence of the IR component of limited type.

Also, if we change the declaration of the Rec record from the previous example and remove the **limited** keyword, the type itself remains implicitly limited. We can see that when trying to assign to objects of Rec type in the Show_Implicitly_Limited procedure:

Listing 70: simple_recs.ads

```
package Simple_Recs is
1
2
       type Int Rec is limited record
3
          V : Integer;
4
      end record;
5
6
      type Rec is record
7
          IR : Int_Rec;
8
      end record;
9
10
   end Simple_Recs;
11
```

Listing 71: show_implicitly_limited.adb

```
with Simple_Recs; use Simple_Recs;
procedure Show_Implicitly_Limited is
A, B : Rec;
begin
B := A;
end Show_Implicitly_Limited;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Record_Components_
    Limited_Type.Record_Components_Limited_Type
MD5: 39770daecfc4579407a799e14f9feff9
```

Build output

Here, the compiler indicates that the assignment is forbidden because the Rec type has a component of limited type. The rationale for this rule is that an object of a limited type doesn't allow assignment or equality, including the case in which that object is a component of some enclosing composite object. If we allowed the enclosing object to be copied or tested for equality, we'd be doing it for all the components, too.

In the Ada Reference Manual

3.8 Record Types³²⁴

17.9 Limited types and aggregates

\rm 1 Note

This section was originally written by Robert A. Duff and published as Gem #1: Limited Types in Ada 2005^{325} and Gem $#2^{326}$.

In this section, we focus on using aggregates to initialize limited types.

• Historically

Prior to Ada 2005, aggregates were illegal for limited types. Therefore, we would be faced with a difficult choice: Make the type limited, and initialize it like this:

Listing 72: persons.ads

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   package Persons is
4
5
      type Limited Person;
6
      type Limited Person Access is
7
        access all Limited Person;
8
9
      type Limited Person is limited record
10
         Name
                    : Unbounded String;
11
          Age
                     : Natural;
12
      end record;
13
14
   end Persons;
15
```

324 http://www.ada-auth.org/standards/22rm/html/RM-3-8.html

³²⁵ https://www.adacore.com/gems/gem-1

³²⁶ https://www.adacore.com/gems/gem-2

```
1
2
3
4
5
6
7
8
9
10
11
```

Listing 73: show_non_aggregate_init.adb

```
with Persons; use Persons;
procedure Show_Non_Aggregate_Init is
X : Limited_Person;
begin
X.Name := To_Unbounded_String ("John Doe");
X.Age := 25;
end Show_Non_Aggregate_Init;
```

Code block metadata

with Ada.Strings.Unbounded;

use Ada.Strings.Unbounded;

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Types_

⇔Aggregates.Full_Coverage_Rules_Limited_Ada95

MD5: fd3dcb6251f7b6912dafcca052932be2
```

which has the maintenance problem the full coverage rules are supposed to prevent. Or, make the type nonlimited, and gain the benefits of aggregates, but lose the ability to prevent copies.

17.9.1 Full coverage rules for limited types

Previously, we discussed *full coverage rules for aggregates* (page 260). They also apply to limited types.

Historically

The full coverage rules have been aiding maintenance since Ada 83. However, prior to Ada 2005, we couldn't use them for limited types.

Suppose we have the following limited type:

Listing 74: persons.ads

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   package Persons is
4
5
      type Limited_Person;
6
      type Limited_Person_Access is
7
        access all Limited Person;
8
9
      type Limited_Person is limited record
10
         Self : Limited_Person_Access :=
11
                   Limited Person'Unchecked Access;
12
         Name : Unbounded_String;
13
         Age : Natural;
14
          Shoe Size : Positive;
15
      end record;
16
17
18
   end Persons;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Types_

⊶Aggregates.Full_Coverage_Rules_Limited

MD5: b8ece44a10d512061cb138be21e42034
```

This type has a self-reference; it doesn't make sense to copy objects, because Self would end up pointing to the wrong place. Therefore, we would like to make the type limited, to prevent developers from accidentally making copies. After all, the type is probably private, so developers using this package might not be aware of the problem. We could also solve that problem with controlled types, but controlled types are expensive, and add unnecessary complexity if not needed.

We can initialize objects of limited type with an aggregate. Here, we can say:

Listing 75: show_aggregate_box_init.adb

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
   with Persons; use Persons;
4
5
   procedure Show_Aggregate_Box_Init is
6
      X : aliased Limited_Person :=
7
             (Self
                         => <>.
8
              Name
                          =>
9
                To Unbounded String ("John Doe"),
10
              Age
                         => 25,
11
              Shoe Size \Rightarrow 10);
12
   begin
13
      null;
14
   end Show Aggregate Box Init;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Types_

⇔Aggregates.Full_Coverage_Rules_Limited

MD5: ded40ff29b53ea5528efba94efaadbec
```

The Self => <> means use the default value of Limited_Person 'Unchecked_Access. Since Limited_Person appears inside the type declaration, it refers to the "current instance" of the type, which in this case is X. Thus, we are setting X.Self to be X'Unchecked_Access.

One very important requirement should be noted: the implementation is required to build the value of X *in place*; it cannot construct the aggregate in a temporary variable and then copy it into X, because that would violate the whole point of limited objects — you can't copy them.

Historically					
Since Ada 2005, an aggregate is allowed to be limited; we can say:					
Listing 76: show_aggregate_init.adb					
with Ada.Strings.Unbounded;					
use Ada.Strings.Unbounded;					
with Persons; use Persons;					
procedure Show Aggregate Init is					
procedure Snow_Aggregate_Init is					
X : aliased Limited Person :=					
(Self => null, Wrong!					
Name =>					
To_Unbounded_String ("John Doe"),					

11 12

```
=> 25,
Aae
Shoe Size \Rightarrow 10);
```

13 14

15

begin

```
X.Self := X'Unchecked Access;
```

end Show_Aggregate_Init;

Code block metadata

```
Project: Courses.Advanced Ada.Resource Management.Limited Types.Limited Types
 →Aggregates.Full Coverage Rules Limited
MD5: 793ee000fd777d0aa5c15e16132ec411
```

It seems uncomfortable to set the value of Self to the wrong value (null) and then correct it. It also seems annoying that we have a (correct) default value for Self, but prior to Ada 2005, we couldn't use defaults with aggregates. Since Ada 2005, a new syntax in aggregates is available: <> means "use the default value, if any". Therefore, we can replace Self => **null** by Self => <>.

Important

Note that using <> in an aggregate can be dangerous, because it can leave some components uninitialized. <> means "use the default value". If the type of a component is scalar, and there is no record-component default, then there is no default value.

For example, if we have an aggregate of type **String**, like this:

Uninitialized Const Str : constant String :=

Listing 77: show string box init.adb

 $(1 \dots 10 \implies <>);$

```
1
2
3
```

4

5

6

begin

null; end Show String Box Init;

procedure Show String Box Init is

Code block metadata

Project: Courses.Advanced Ada.Resource Management.Limited Types.Limited Types ⇔Aggregates.String Box Init MD5: 28931ced4e1113d55bdc9dc64b42f70a

we end up with a 10-character string all of whose characters are invalid values. Note that this is no more nor less dangerous than this:

Listing 78: show dangerous string.adb

```
procedure Show_Dangerous_String is
1
       Uninitialized_String_Var : String (1 .. 10);
2
3
       -- no initialization
4
5
       Uninitialized Const Str : constant String :=
6
           Uninitialized_String_Var;
7
   begin
8
      null;
9
10
```

end Show_Dangerous_String;

Code block metadata

```
Project: Courses.Advanced Ada.Resource Management.Limited Types.Limited Types
→Aggregates.Dangerous_String
MD5: 6c26e9c8d5d031d4e6eac1ac8458f17e
```

Build output

As always, one must be careful about uninitialized scalar objects.

17.10 Constructor functions for limited types

\rm 1 Note

This section was originally written by Robert A. Duff and published as Gem #3³²⁷.

Given that we can use build-in-place aggregates for limited types, the obvious next step is to allow such aggregates to be wrapped in an abstraction — namely, to return them from functions. After all, interesting types are usually private, and we need some way for clients to create and initialize objects.

Historically

Prior to Ada 2005, constructor functions (that is, functions that create new objects and return them) were not allowed for limited types. Since Ada 2005, fully-general constructor functions are allowed.

Let's see an example:

```
Listing 79: p.ads
```

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   package P is
4
      task type Some_Task_Type;
5
6
      protected type Some Protected Type is
7
         -- dummy type
8
      end Some Protected Type;
9
10
       type T (<>) is limited private;
11
       function Make T (Name : String) return T;
12
                ~~~~~
13
       - -
       -- constructor function
14
   private
15
       type T is limited
16
         record
17
                    : Unbounded String:
             Name
18
             My_Task : Some_Task Type;
19
             My Prot : Some Protected Type;
20
          end record;
21
   end P;
22
```

³²⁷ https://www.adacore.com/gems/gem-3

```
Listing 80: p.adb
```

```
package body P is
1
2
      task body Some Task Type is
3
      begin
4
          null;
5
      end Some_Task_Type;
6
7
      protected body Some_Protected_Type is
8
      end Some_Protected_Type;
9
10
      function Make_T (Name : String) return T is
11
      begin
12
          return (Name
                        =>
13
                     To_Unbounded_String (Name),
14
                  others => <>);
15
      end Make T;
16
17
   end P;
18
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Constructor_

⊶Functions_Limited_Types.Constructor_Functions

MD5: 2e73eea0ba7852d45ba96dc1f6fae14d
```

Given the above, clients can say:

Listing 81: show_constructor_function.adb

```
with P; use P;
procedure Show_Constructor_Function is
My_T : T := Make_T
(Name => "Bartholomew Cubbins");
begin
null;
end Show_Constructor_Function;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Constructor_

⊶Functions_Limited_Types.Constructor_Functions

MD5: 52801fafbd58fedbf268a6704008627b
```

As for aggregates, the result of Make_T is built in place (that is, in My_T), rather than being created and then copied into My_T. Adding another level of function call, we can do:

Listing 82: show_rumplestiltskin_constructor.adb

```
with P; use P;
1
2
   procedure Show_Rumplestiltskin_Constructor is
3
4
      function Make_Rumplestiltskin return T is
5
      begin
6
           return Make_T (Name => "Rumplestiltskin");
7
      end Make_Rumplestiltskin;
8
9
      Rumplestiltskin_Is_My_Name : constant T :=
10
```

```
Make_Rumplestiltskin;
11
```

- begin 12 null;
- 13

```
end Show_Rumplestiltskin_Constructor;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Constructor_
→Functions_Limited_Types.Constructor_Functions
MD5: d8d9e9f22a0f2f034057fe97f75eacfe
```

It might help to understand the implementation model: In this case, Rumplestiltskin Is My Name is allocated in the usual way (on the stack, presuming it is declared local to some subprogram). Its address is passed as an extra implicit parameter to Make Rumplestiltskin, which then passes that same address on to Make T, which then builds the aggregate in place at that address. Limited objects must never be copied! In this case, Make_T will initialize the Name component, and create the My Task and My Prot components, all directly in Rumplestiltskin Is My Name.

Historically

Note that Rumplestiltskin Is My Name is constant. Prior to Ada 2005, it was impossible to create a constant limited object, because there was no way to initialize it.

As we discussed before (page 812), the (<>) on type T means that it has unknown discrim*inants* from the point of view of the client. This is a trick that prevents clients from creating default-initialized objects (that is, X : T; is illegal). Thus clients must call Make T whenever an object of type T is created, giving package P full control over initialization of objects.

Ideally, limited and nonlimited types should be just the same, except for the essential difference: you can't copy limited objects (and there's no language-defined equality operator). By allowing functions and aggregates for limited types, we're very close to this goal. Some languages have a specific feature called *constructor*. In Ada, a *constructor* is just a function that creates a new object.

Historically

1 2

3

4

5

6

7

8

9

10

Prior to Ada 2005, constructors only worked for nonlimited types. For limited types, the only way to *construct* on declaration was via default values, which limits you to one constructor. And the only way to pass parameters to that construction was via discriminants.

Consider the following package:

Listing 83: aux.ads

with Ada.Containers.Ordered Sets;

```
package Aux is
   generic
      with package OS is new
        Ada.Containers.Ordered Sets (<>);
   function Gen_Singleton_Set
     (Element : OS.Element_Type)
      return OS.Set:
end Aux:
```

```
Listing 84: aux.adb
   package body Aux is
1
      function Gen_Singleton_Set
2
         (Element : OS.Element Type)
3
         return OS.Set
4
      is
5
      begin
6
         return S : OS.Set := OS.Empty Set do
7
            S.Insert (Element);
8
         end return;
9
      end Gen_Singleton_Set;
10
   end Aux;
11
   Code block metadata
   Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Constructor_
    Generations_Limited_Types.Constructor_Functions_2
   MD5: b715ae504c49ed59b7fd5ead4cc7bbb4
   Since Ada 2005, we can say:
                               Listing 85: show set decl.adb
   with Ada.Containers.Ordered_Sets;
1
   with Aux;
2
3
   procedure Show_Set_Decl is
4
5
      package Integer_Sets is new
6
        Ada.Containers.Ordered Sets
7
          (Element_Type => Integer);
8
      use Integer_Sets;
9
10
      function Singleton_Set is new
11
        Aux.Gen_Singleton_Set
12
           (0S => Integer_Sets);
13
14
      This Set : Set := Empty Set;
15
      That_Set : Set := Singleton_Set
16
17
                            (Element => 42);
   begin
18
      null;
19
   end Show_Set_Decl;
20
   Code block metadata
   Project: Courses.Advanced Ada.Resource Management.Limited Types.Constructor
    →Functions Limited Types.Constructor Functions 2
   MD5: 443fc3390b0f3e5516d91c80f16bed3f
   whether or not Set is limited. This Set : Set := Empty Set; seems clearer than:
```

```
with Ada.Containers.Ordered Sets;
1
2
   procedure Show_Set_Decl is
3
4
       package Integer Sets is new
5
         Ada.Containers.Ordered Sets
6
           (Element Type => Integer);
7
      use Integer Sets;
8
9
      This_Set : Set;
10
   begin
11
      null;
12
   end Show_Set_Decl;
13
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Constructor_ ⊶Functions_Limited_Types.Constructor_Functions_2 MD5: e5b6c0e148cfdb1987ab3002ec1f53bd

which might mean "default-initialize to the empty set" or might mean "leave it uninitialized, and we'll initialize it in later".

Listing 86: show set decl.adb

17.11 Return objects

17.11.1 Extended return statements for limited types

\rm 1 Note

This section was originally written by Robert A. Duff and published as Gem #10: Limited Types in Ada 2005³²⁸.

Previously, we discussed *extended return statements* (page 462). For most types, extended return statements are no big deal — it's just syntactic sugar. But for limited types, this syntax is almost essential:

```
Listing 87: task_construct_error.ads
```

```
1 package Task_Construct_Error is
2
3 task type Task_Type (Discriminant : Integer);
4
5 function Make_Task (Val : Integer)
6 return Task_Type;
7
8 end Task Construct Error;
```

Listing 88: task construct error.adb

```
1 package body Task_Construct_Error is
2
3 task body Task_Type is
4 begin
5 null;
```

```
<sup>328</sup> https://www.adacore.com/gems/ada-gem-10
```

```
end Task_Type;
6
7
      function Make_Task (Val : Integer)
8
                             return Task_Type
9
10
      is
          Result : Task_Type
11
                      (Discriminant => Val * 3);
12
      begin
13
             some statements...
14
          return Result; -- Illegal!
15
      end Make_Task;
16
17
   end Task_Construct_Error;
18
```

Code block metadata

The return statement here is illegal, because Result is local to Make_Task, and returning it would involve a copy, which makes no sense (which is why task types are limited). Since Ada 2005, we can write constructor functions for task types:

Listing 89: task_construct.ads

```
1 package Task_Construct is
2
3 task type Task_Type (Discriminant : Integer);
4
5 function Make_Task (Val : Integer)
6 return Task_Type;
7
8 end Task Construct;
```

Listing 90: task_construct.adb

```
package body Task Construct is
1
2
       task body Task_Type is
3
4
      begin
5
          null;
      end Task_Type;
6
7
      function Make_Task (Val : Integer)
8
                             return Task_Type is
9
       begin
10
          return Result : Task_Type
11
                              (Discriminant => Val * 3)
12
          do
13
             -- some statements...
14
             null;
15
          end return;
16
17
      end Make_Task;
18
   end Task_Construct;
19
```

Code block metadata

```
Gatements_Limited_Types.Extended_Return_Limited MD5: c91a24f09a76aef1c25d1a55bcbee910
```

If we call it like this:

Listing 91: show task construct.adb

```
with Task_Construct; use Task_Construct;
procedure Show_Task_Construct is
My_Task : Task_Type := Make_Task (Val => 42);
begin
null;
end Show_Task_Construct;
```

Code block metadata

Result is created *in place* in My_Task. Result is temporarily considered local to Make_Task during the -- *some statements* part, but as soon as Make_Task returns, the task becomes more global. Result and My_Task really are one and the same object.

When returning a task from a function, it is activated after the function returns. The -- some statements part had better not try to call one of the task's entries, because that would deadlock. That is, the entry call would wait until the task reaches an accept statement, which will never happen, because the task will never be activated.

17.11.2 Initialization and function return

As mentioned in the previous section, the object of limited type returned by the initialization function is built *in place*. In other words, the return object is built in the object that is the target of the assignment statement.

For example, we can see this when looking at the address of the object *returned* by the Init function, which we call to initialize the limited type Simple_Rec:

Listing 92: limited_types.ads

```
package Limited Types is
1
2
      type Integer_Access is access Integer;
3
4
      type Simple_Rec is limited private;
5
6
      function Init (I : Integer) return Simple_Rec;
7
8
   private
9
10
      type Simple_Rec is limited record
11
         V : Integer Access;
12
      end record;
13
14
   end Limited_Types;
15
```

Listing 93: limited_types.adb

```
with Ada.Text_I0;
                                  use Ada.Text_I0;
1
   with System;
2
   with System.Address_Image;
3
4
   package body Limited_Types is
5
6
      function Init (I : Integer) return Simple_Rec
7
      is
8
      begin
9
          return E : Simple_Rec do
10
             E.V := new Integer'(I);
11
12
             Put Line ("E'Address (Init):
13
                        & System.Address_Image
14
                             (E'Address));
15
          end return;
16
      end Init;
17
18
   end Limited_Types;
19
```

Listing 94: show_limited_init.adb

```
with Ada.Text_I0;
                                  use Ada.Text_I0;
1
   with System;
2
3
   with System.Address_Image;
4
   with Limited_Types;
                                  use Limited_Types;
5
6
   procedure Show_Limited_Init is
7
   begin
8
      declare
9
          A : Simple_Rec := Init (0);
10
      begin
11
          Put_Line ("A'Address (local): "
12
                     & System.Address_Image
13
                         (A'Address));
14
      end:
15
      Put_Line ("----");
16
17
      declare
18
          B : Simple Rec := Init (0);
19
      begin
20
          Put Line ("B'Address (local): "
21
                     & System.Address Image
22
                         (B'Address));
23
24
       end;
   end Show_Limited_Init;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Extended_Return_

→Statements_Limited_Types.Initialization_Return_Do

MD5: 67235f804206e07fa4eba3a45cc1096f
```

Runtime output

```
E'Address (Init): 00007FFCF1C56F38
A'Address (local): 00007FFCF1C56F38
```

```
E'Address (Init): 00007FFCF1C56F30
B'Address (local): 00007FFCF1C56F30
```

When running this code example and comparing the address of the object E in the Init function and the object that is being initialized in the Show_Limited_Init procedure, we see that the return object E (of the Init function) and the local object in the Show_Limited_Init procedure are the same object.

```
Important
When we use nonlimited types, we're actually copying the returned object — which was
locally created in the function — to the object that we're assigning the function to.
For example, let's modify the previous code and make Simple Rec nonlimited:
                          Listing 95: non limited types.ads
         package Non_Limited_Types is
      1
      2
            type Integer_Access is access Integer;
      3
      4
            type Simple_Rec is private;
      5
      6
            function Init (I : Integer)
      7
                            return Simple_Rec;
      8
      9
         private
     10
     11
            type Simple Rec is record
     12
               V : Integer_Access;
     13
            end record;
     14
     15
     16 end Non_Limited_Types;
                          Listing 96: non limited types.adb
         with Ada.Text_I0;
                                      use Ada.Text_I0;
      1
         with System;
      2
         with System.Address_Image;
      3
      4
         package body Non_Limited_Types is
      5
      6
            function Init (I : Integer)
      7
                            return Simple Rec is
      8
            begin
      9
               return E : Simple_Rec do
     10
                  E.V := new Integer'(I);
     11
     12
                  Put_Line ("E'Address (Init):
     13
                             & System.Address_Image
     14
                                 (E'Address));
     15
               end return;
     16
            end Init;
     17
     18
        end Non Limited Types;
     19
```

```
Listing 97: show non limited init by copy.adb
          with Ada.Text I0;
                                                                                                          use Ada.Text I0;
  1
           with System;
 2
           with System.Address_Image;
  3
  4
           with Non Limited Types;
  5
          use Non Limited Types;
  6
           procedure Show Non Limited Init By Copy is
 8
                   A, B : Simple Rec;
 9
          begin
10
                     declare
11
                              A : Simple Rec := Init (0);
12
                     beain
13
                              Put_Line ("A'Address (local): "
14
                                                                & System.Address Image
15
                                                                               (A'Address)):
16
                     end;
17
                     Put Line ("----");
18
19
20
                     declare
                          B : Simple_Rec := Init (0);
21
                     begin
22
                               Put_Line ("B'Address (local): "
23
                                                                 & System.Address_Image
24
                                                                               (B'Address));
25
                      end:
26
        end Show Non Limited Init By Copy;
27
            Code block metadata
           Project: Courses.Advanced Ada.Resource Management.Limited
               Graphic Statements St
                →Initialization Return Copy
           MD5: 6e224b64b90dabdf5064c70364fa80cb
           Runtime output
           E'Address (Init): 00007FFFBE241EB0
           A'Address (local): 00007FFFBE241FA8
```

E'Address (Init): 00007FFFBE241EB0 B'Address (local): 00007FFFBE241FA0

In this case, we see that the local object E in the Init function is not the same as the object it's being assigned to in the Show_Non_Limited_Init_By_Copy procedure. In fact, E is being copied to A and B.

17.12 Building objects from constructors

\rm Note

This section was originally written by Robert A. Duff and published as Gem #11: Limited Types in Ada 2005³²⁹.

We've earlier seen examples of constructor functions for limited types similar to this:

```
<sup>329</sup> https://www.adacore.com/gems/ada-gem-11
```

```
Listing 98: p.ads
```

```
with Ada.Strings.Unbounded;
1
   use Ada.Strings.Unbounded;
2
3
   package P is
4
      task type Some_Task_Type;
5
6
      protected type Some Protected Type is
7
          -- dummy type
8
      end Some Protected Type;
9
10
       type T is limited private;
11
      function Make_T (Name : String) return T;
12
                ~~~~~
13
       - -
       -- constructor function
14
   private
15
       type T is limited
16
          record
17
             Name
                     : Unbounded String;
18
             My_Task : Some_Task_Type;
19
             My_Prot : Some_Protected_Type;
20
          end record;
21
   end P;
22
```

```
Listing 99: p.adb
```

```
package body P is
1
2
      task body Some_Task_Type is
3
      begin
4
          null;
5
      end Some_Task_Type;
6
7
      protected body Some_Protected_Type is
8
      end Some Protected Type;
9
10
      function Make_T (Name : String) return T is
11
      begin
12
          return (Name
13
                         =>
                    To_Unbounded_String (Name),
14
                  others => <>);
15
      end Make_T;
16
17
   end P;
18
```

Listing 100: p-aux.ads

```
1 package P.Aux is
2 function Make_Rumplestiltskin return T;
3 end P.Aux;
```

Listing 101: p-aux.adb

```
1 package body P.Aux is
2
3 function Make_Rumplestiltskin return T is
4 begin
5 return Make_T (Name => "Rumplestiltskin");
6 end Make_Rumplestiltskin;
```

7

8 end P.Aux;

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Building_Objects_

⊶From_Constructors.Building_Objs_From_Constructors

MD5: 1956721292a82899d244afcd10ff63ed
```

It is useful to consider the various contexts in which these functions may be called. We've already seen things like:

Listing 102: show rumplestiltskin constructor.adb

```
with P;
               use P;
1
  with P.Aux; use P.Aux;
2
3
  procedure Show Rumplestiltskin Constructor is
4
      Rumplestiltskin Is My Name : constant T :=
5
        Make Rumplestiltskin;
6
7
  begin
8
      null:
  end Show Rumplestiltskin Constructor;
9
```

Code block metadata

in which case the limited object is built directly in a standalone object. This object will be finalized whenever the surrounding scope is left.

We can also do:

```
Listing 103: show_parameter_constructor.adb
```

```
1 with P; use P;
2 with P.Aux; use P.Aux;
3
4 procedure Show_Parameter_Constructor is
5 procedure Do_Something (X : T) is null;
6 begin
7 Do_Something (X => Make_Rumplestiltskin);
8 end Show_Parameter_Constructor;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Building_Objects_

⊶From_Constructors.Building_Objs_From_Constructors

MD5: 6lccaefb4b7cfc42c065aa15543fc13b
```

Here, the result of the function is built directly in the formal parameter X of Do_Something. X will be finalized as soon as we return from Do_Something.

We can allocate initialized objects on the heap:

Listing 104: show_heap_constructor.adb

```
with P; use P;
```

```
with P.Aux; use P.Aux;
```

```
3
   procedure Show_Heap_Constructor is
4
5
       type T_Ref is access all T;
6
7
      Global : T_Ref;
8
9
      procedure Heap_Alloc is
10
          Local : T_Ref;
11
          To_Global : Boolean := True;
12
      begin
13
          Local := new T'(Make_Rumplestiltskin);
14
          if To_Global then
15
             Global := Local;
16
17
          end if;
      end Heap_Alloc;
18
19
   begin
20
      null:
21
   end Show_Heap_Constructor;
22
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Building_Objects_

⊶From_Constructors.Building_Objs_From_Constructors

MD5: 8eb794884f1dfbdbedf1bc4369f45cf8
```

The result of the function is built directly in the heap-allocated object, which will be finalized when the scope of T_Ref is left (long after Heap_Alloc returns).

We can create another limited type with a component of type T, and use an aggregate:

Listing 105: show_outer_type.adb

```
use P;
   with P;
1
   with P.Aux; use P.Aux;
2
3
   procedure Show_Outer_Type is
4
5
       type Outer_Type is limited record
6
          This : T;
7
          That : T;
8
      end record;
9
10
      Outer_Obj : Outer_Type :=
11
                     (This => Make_Rumplestiltskin,
12
                      That => Make_T (Name => ""));
13
14
   begin
15
      null;
16
   end Show_Outer_Type;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Building_Objects_

⊶From_Constructors.Building_Objs_From_Constructors

MD5: 00817649406492b79977d67eb0fd3955
```

As usual, the function results are built in place, directly in Outer_Obj.This and Outer_Obj. That, with no copying involved.

The one case where we cannot call such constructor functions is in an assignment state-

ment:

Listing 106: show_illegal_constructor.adb

```
with P;
               use P;
1
  with P.Aux; use P.Aux;
2
3
  procedure Show Illegal Constructor is
4
      Rumplestiltskin_Is_My_Name : T;
5
  begin
6
      Rumplestiltskin_Is_My_Name :=
7
        Make_T (Name => ""); -- Illegal!
8
  end Show Illegal Constructor;
a
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Building_Objects_ Gerom_Constructors.Building_Objs_From_Constructors MD5: f7b0c78e9fbe2e104b82dfff25ac3e3a

Build output

```
show_illegal_constructor.adb:7:04: error: left hand of assignment must not be_
__limited type
gprbuild: *** compilation phase failed
```

which is illegal because assignment statements involve copying. Likewise, we can't copy a limited object into some other object:

Listing 107: show_illegal_constructor.adb

```
with P;
                use P;
1
   with P.Aux; use P.Aux;
2
3
   procedure Show Illegal Constructor is
4
      Rumplestiltskin Is My Name : constant T :=
5
        Make_T (Name => "");
6
      Other : T :=
7
        Rumplestiltskin_Is_My_Name; -- Illegal!
8
   begin
9
      null:
10
  end Show_Illegal_Constructor;
11
```

17.13 Limited types as parameter

Previously, we saw that *parameters can be passed by copy or by reference* (page 465). Also, we discussed the concept of by-copy and by-reference types. *Explicitly limited types* (page 795) are by-reference types. Consequently, parameters of these types are always passed by reference.

I For further reading...

As an example of the importance of this rule, consider the case of a lock (as an abstract data type). If such a lock object were passed by copy, the Acquire and Release operations would be working on copies of this object, not on the original one. This would lead to timing-dependent bugs.

Let's reuse an example of an explicitly limited type:

Listing 108: simple_recs.ads

```
1 package Simple_Recs is
2
3 type Rec is limited record
4 I : Integer;
5 end record;
6
7 end Simple_Recs;
```

Code block metadata

In this example, Rec is a by-reference type because the type declaration is an explicit limited record. Therefore, the parameter R of the Proc procedure is passed by reference.

We can run the Test application below and compare the address of the R object from Test to the address of the R parameter of Proc to determine whether both R s refer to the same object or not:

Listing 109: simple_recs.ads

```
with System;
1
2
   package Simple Recs is
3
4
      type Rec is limited record
5
6
         I : Integer;
7
      end record;
8
      procedure Proc (R : in out Rec;
9
                       A : out System.Address);
10
11
   end Simple_Recs;
12
```

Listing 110: simple_recs.adb

```
package body Simple_Recs is
1
2
      procedure Proc (R : in out Rec;
3
                      A : out System.Address) is
4
      begin
5
         R.I := 0;
6
         A := R'Address;
7
      end Proc;
8
9
   end Simple_Recs;
10
```

Listing 111: test.adb

```
with Ada.Text_I0;
                                use Ada.Text_I0;
1
  with System;
                                use System;
2
  with System.Address Image;
3
  with Simple_Recs;
                                use Simple_Recs;
4
5
  procedure Test is
6
      R : Rec:
7
8
```

```
AR_Proc, AR_Test : System.Address;
9
10
   beain
      AR Proc := R'Address;
11
12
      Proc (R, AR_Test);
13
14
      Put_Line ("R'Address (Proc): "
15
                 & System.Address_Image (AR_Proc));
16
      Put_Line ("R'Address (Test):
17
                 & System.Address_Image (AR_Test));
18
19
       if AR_Proc = AR_Test then
20
          Put_Line ("R was passed by reference.");
21
22
       else
          Put_Line ("R was passed by copy.");
23
       end if;
24
25
   end Test;
26
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Limited_Types.Limited_Types_ ⊲Parameters.Explicitly_Limited_Types MD5: d4fe2bb47d2223ef013d22aa305403e5

Runtime output

R'Address (Proc): 00007FFEBCAF32EC R'Address (Test): 00007FFEBCAF32EC R was passed by reference.

When running the Test application, we confirm that R was passed by reference. Note, however, that the fact that R was passed by reference doesn't automatically imply that Rec is a by-reference type: the type could have been ambiguous, and the compiler could have just decided to pass the parameter by reference in this case.

Therefore, we have to rely on the rules specified in the Ada Reference Manual:

- 1. If a limited type is explicitly limited, a parameter of this type is a by-reference type.
 - The rule applies to all kinds of explicitly limited types. For example, consider private limited types where the type is declared limited in the private type's completion (in the package's private part): a parameter of this type is a by-reference type.
- 2. If a limited type is not *explicitly* limited, a parameter of this type is neither a by-copy nor a by-reference type.
 - In this case, the decision whether the parameter is passed by reference or by copy is made by the compiler.

In the Ada Reference Manual

- 6.2 Formal Parameter Modes³³⁰
- 6.4.1 Parameter Associations³³¹
- 7.5 Limited Types³³²

³³⁰ http://www.ada-auth.org/standards/22rm/html/RM-6-2.html

³³¹ http://www.ada-auth.org/standards/22rm/html/RM-6-4-1.html

³³² http://www.ada-auth.org/standards/22rm/html/RM-7-5.html

CHAPTER EIGHTEEN

CONTROLLED TYPES

18.1 Overview

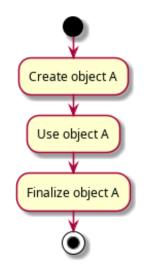
In this section, we introduce the concept of controlled types. We start with a review of lifetime of objects and discuss how controlled types allow us to control the initialization, post-copy (e.g. assignment) adjustment and finalization of objects.

1 Relevant topics

• Assignment and Finalization³³³

18.1.1 Lifetime of objects

We already talked about the lifetime of objects³³⁴ previously in the context of *access types* (page 645). Again, we assume you understand the concept. In any case, let's quickly review the typical lifetime of an object:



In simple terms, an object A is first created before we can make use of it. When object A is about to get out of scope, it is finalized. Note that finalization might not entail any actual code execution — but it often does.

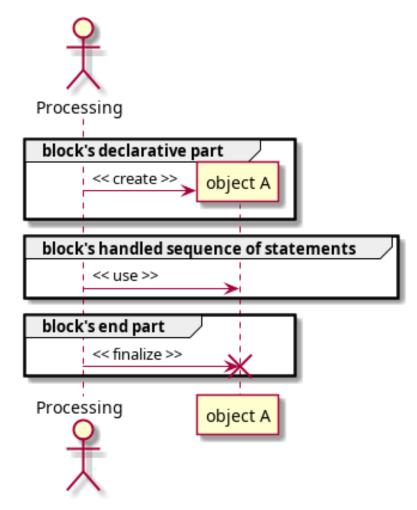
Let's analyze the lifetime of object A in a procedure P:

³³³ http://www.ada-auth.org/standards/22rm/html/RM-7-6.html

³³⁴ https://en.wikipedia.org/wiki/Variable_(computer_science)#Scope_and_extent

procedure P is
 A : T;
begin
 P2 (A);
end P;

We could visualize the lifetime as follows:



In other words, object A is created in the declarative part of P and then it's used in P's sequence of statements. Finally, A is finalized when P ends.

18.1.2 Initialization of objects

Typically, right after an object A is created, it is still uninitialized. Therefore, we have to explicitly initialize it with a meaningful initial value — or with the value returned by a function call, for example. Similarly, when an object A is about to get out of scope, it is going to be finalized (i.e. destroyed) and its contents are then lost forever.

As we know, for some standard Ada types, objects are initialized by default. For example, objects of access types are initialized by default to **null**. Likewise, we can declare *types* with default initial value (page 65):

Listing 1: main.adb

```
with Ada.Text_I0; use Ada.Text_I0;
```

2

```
procedure Main is
3
      type Int is new Integer
4
        with Default_Value => 42;
5
6
      I : Int;
7
      AI : access Int;
8
   begin
9
      Put_Line ("I : "
10
                 & I'Image);
11
      Put_Line ("AI : "
12
                 & AI'Image);
13
   end Main;
14
```

Code block metadata

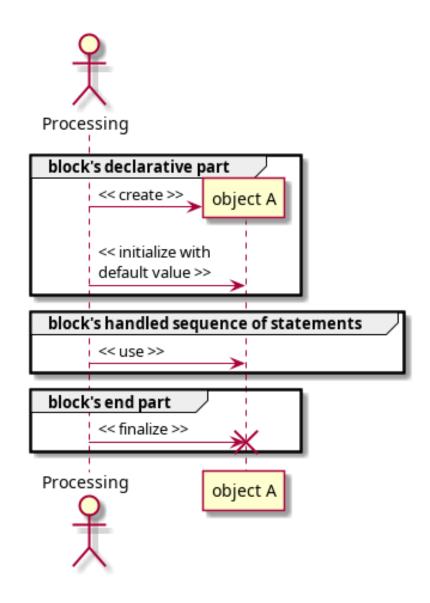
Build output

main.adb:8:04: warning: variable "AI" is read but never assigned [-gnatwv]

Runtime output

I: 42 AI:null

In this case, we can visualize the lifetime of those objects as follows:



Even though these default initialization methods provide some control over the objects, they might not be enough in certain situations. Also, we don't have any means to perform useful operations right before an object gets out of scope.

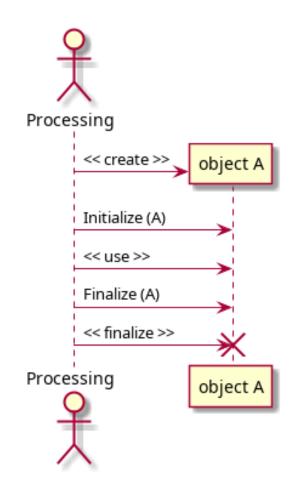
• For further reading...

In general, record types have a very good default initialization capability. They're the most common completion for private types, so the facility is often used. In this sense, default initialization is the first choice, as it's guaranteed and requires nothing of the client. In addition, it's cheap at run-time compared to controlled types.

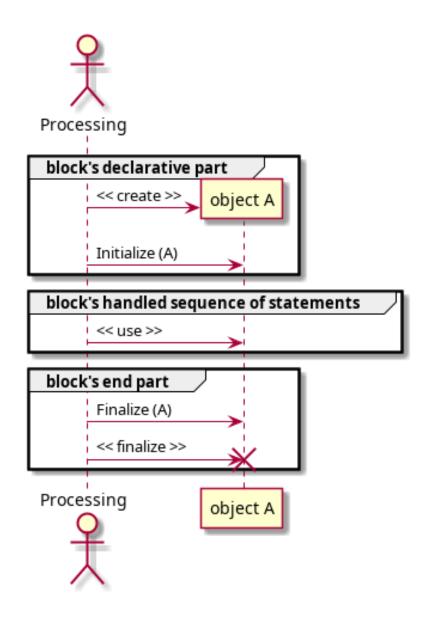
18.1.3 Controlled objects

Controlled objects allow us to better control the initialization and finalization of an object. For any controlled object A, an Initialize (A) procedure is called right *after* the object is created, and a Finalize (A) procedure is called right *before* the object is actually finalized.

We can visualize the lifetime of controlled objects as follows:



In the context of a block statement, the lifetime becomes:



Let's look at a simple example:



```
with Ada.Finalization;
1
2
   package Simple_Controlled_Types is
3
4
      type T is tagged private;
5
6
      procedure Dummy (E : T);
7
8
   private
9
10
      type T is new
11
        Ada.Finalization.Controlled
12
          with null record;
13
14
      overriding
15
      procedure Initialize (E : in out T);
16
17
      overriding
18
```

```
procedure Finalize (E : in out T);
end Simple_Controlled_Types;
```

Listing 3: simple_controlled_types.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   package body Simple_Controlled_Types is
3
4
      procedure Dummy (E : T) is
5
      begin
6
          Put_Line ("(Dummy...)");
7
      end Dummy;
8
9
      procedure Initialize (E : in out T) is
10
      begin
11
          Put Line ("Initialize...");
12
      end Initialize;
13
14
      procedure Finalize (E : in out T) is
15
      begin
16
          Put Line ("Finalize...");
17
      end Finalize;
18
19
   end Simple_Controlled_Types;
20
```

Listing 4: show controlled types.adb

```
with Simple Controlled Types;
1
   use Simple_Controlled_Types;
2
3
   procedure Show_Controlled_Types is
4
5
       A : T;
6
       - -
       - -
          This declaration roughly
7
8
       - -
          corresponds to:
9
       - -
       - -
              A : T;
10
          begin
       - -
11
       - -
              Initialize (A);
12
       - -
13
   begin
14
15
       Dummy (A);
16
       -- When A is about to get out of
17
       - -
          scope:
18
19
       - -
       -- Finalize (A);
20
21
   end Show_Controlled_Types;
22
```

Code block metadata

Runtime output

Initialize... (Dummy...) Finalize...

When we run this application, we see the user messages indicating the calls to Initialize and Finalize.

For further reading...

Note that if a controlled object isn't used in the application, the compiler might optimize it out. In this case, procedures Initialize and Finalize won't be called for this object, as it doesn't actually exist. You can see this effect by replacing the call to Dummy (A) in the Show_Controlled_Types procedure by a null statement (**null**).

18.1.4 Adjustment of controlled objects

An assignment is a full bit-wise copy of the entire right-hand side to the entire left-hand side. When copying controlled objects, however, we might need to adjust the target object. This is made possible by overriding the Adjust procedure, which is called right after the copy to an object has been performed. (As we'll see later on, *limited controlled types* (page 844) do not offer an Adjust procedure.)

The deep $copy^{335}$ of objects is a typical example where adjustments are necessary. When we assign an object B to an object A, we're essentially doing a shallow $copy^{336}$. If we have references to other objects in the source object B, those references will be copied as well, so both target A and source B will be referring to the same objects. When performing a deep copy, however, we want the information from the dereferenced objects to be copied, not the references themselves. Therefore, we have to first allocate new objects for the target object A and copy the information from the original references — the ones we copied from the source object B — to the new objects. This kind of processing can be performed in the Adjust procedure.

As an example, let's extend the previous code example and override the Adjust procedure:

Listing 5: simple_controlled_types.a	ads
--------------------------------------	-----

```
with Ada.Finalization;
1
2
   package Simple Controlled Types is
3
4
      type T is tagged private;
5
6
      procedure Dummy (E : T);
7
8
   private
9
10
       type T is new
11
         Ada.Finalization.Controlled
12
           with null record;
13
14
       overriding
15
      procedure Initialize (E : in out T);
16
17
       overriding
18
       procedure Adjust (E : in out T);
19
20
```

(continues on next page)

³³⁵ https://en.wikipedia.org/wiki/Object_copying#Deep_copy
 ³³⁶ https://en.wikipedia.org/wiki/Object_copying#Shallow_copy

```
21 overriding
22 procedure Finalize (E : in out T);
23
24 end Simple_Controlled_Types;
```

Listing 6: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Controlled_Types is
3
4
      procedure Dummy (E : T) is
5
      begin
6
          Put_Line ("(Dummy...)");
7
      end Dummy;
8
9
      procedure Initialize (E : in out T) is
10
11
      begin
          Put_Line ("Initialize...");
12
      end Initialize;
13
14
      procedure Adjust (E : in out T) is
15
      begin
16
          Put_Line ("Adjust...");
17
      end Adjust;
18
19
      procedure Finalize (E : in out T) is
20
      begin
21
          Put_Line ("Finalize...");
22
       end Finalize;
23
24
   end Simple_Controlled_Types;
25
```

Listing 7: show_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Simple_Controlled_Types;
3
   use Simple_Controlled_Types;
4
5
   procedure Show_Controlled_Types is
6
      A, B : T;
7
   begin
8
9
      Put_Line ("A := B");
      A := B;
10
11
      Dummy (A);
12
      Dummy (B);
13
   end Show_Controlled_Types;
14
```

Code block metadata

Runtime output

```
Initialize...
Initialize...
```

A := B Finalize... Adjust... (Dummy...) (Dummy...) Finalize... Finalize...

When running this application, we see that the Adjust procedure is called for object A — right after B is copied to A as part of the A := B assignment. We discuss more about this procedure *later on* (page 858).

18.1.5 Limited controlled types

Ada offers controlled types in two flavors: nonlimited controlled types — such as the ones we've seen so far — and limited controlled types. Both types are declared in the Ada. Finalization package.

The only difference between these types is that limited controlled types don't have an Adjust procedure that could be overridden, as limited types *do not permit direct copies of objects to be made via assignments* (page 782). (Obviously, both controlled and limited controlled types provide Initialize and Finalize procedures.)

The following table summarizes the information:

Туре	Name	Initialize	Finalize	Adjust
Nonlimited Controlled	Controlled	Yes	Yes	Yes
Limited controlled	Limited_Controlled	Yes	Yes	Not available

18.1.6 Simple Example with ID

Although the previous code examples indicated that Initialize, Finalize and Adjust are called as we expect for controlled objects, they didn't show us exactly how those objects are actually handled. In this section, we discuss this by analyzing a code example that assigns a unique ID to each controlled object.

Let's start with the complete code example:

```
Listing 8: simple_controlled_types.ads
```

```
with Ada.Finalization:
1
2
   package Simple Controlled Types is
3
4
      type T is tagged private;
5
6
      procedure Show (E
                            : T:
7
                        Name : String);
8
9
   private
10
11
      protected Id Gen is
12
          procedure New_Id (Id_Out : out Positive);
13
      private
14
         Id : Natural := 0;
15
      end Id_Gen;
16
17
      type T is new
18
```

```
Ada.Finalization.Controlled with
19
       record
20
          Id : Positive;
21
       end record;
22
23
       overriding
24
       procedure Initialize (E : in out T);
25
26
       overriding
27
       procedure Adjust (E : in out T);
28
29
       overriding
30
       procedure Finalize (E : in out T);
31
32
   end Simple_Controlled_Types;
33
```

Listing 9: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Controlled_Types is
3
4
      protected body Id_Gen is
5
6
          procedure New_Id (Id_Out : out Positive) is
7
          begin
8
             Id := Id + 1;
9
             Id Out := Id;
10
11
          end New_Id;
12
      end Id_Gen;
13
14
      procedure Initialize (E : in out T) is
15
       begin
16
          Id_Gen.New_Id (E.Id);
17
          Put_Line ("Initialize: ID => "
18
                     & E.Id'Image);
19
       end Initialize;
20
21
      procedure Adjust (E : in out T) is
22
23
          Prev_Id : constant Positive := E.Id;
24
      begin
          Id_Gen.New_Id (E.Id);
25
                                   ID => "
          Put_Line ("Adjust:
26
                     & E.Id'Image);
27
          Put_Line ("
                         (Previous ID => "
28
                     & Prev Id'Image
29
                     & ")");
30
      end Adjust;
31
32
      procedure Finalize (E : in out T) is
33
      begin
34
          Put_Line ("Finalize:
                                  ID => "
35
                     & E.Id'Image);
36
      end Finalize;
37
38
      procedure Show (E
                            : T:
39
                        Name : String) is
40
      begin
41
          Put_Line ("Obj. " & Name
42
                     & ": ID => "
43
```

44 & & E.Id'Image);
45 end Show;
46
47 end Simple_Controlled_Types;

Listing 10: show_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Simple_Controlled_Types;
3
   use Simple_Controlled_Types;
4
5
   procedure Show_Controlled_Types is
6
       A, B : T;
7
8
       -- Declaration corresponds to:
9
       - -
10
       -- declare
11
       - -
             A, B : T;
12
       -- begin
13
             Initialize (A);
14
       - -
              Initialize (B);
15
       - -
       -- end;
16
   begin
17
       Put_Line ("-----");
18
       Show (A, "A");
Show (B, "B");
19
20
21
       Put_Line ("-----");
22
       Put_Line ("A := B;");
23
24
       A := B;
25
       -- Statement corresponds to:
26
       - -
27
       -- Finalize (A);
28
       -- A := B;
29
       -- Adjust (A);
30
31
       Put Line ("-----");
32
       Show (A, "A");
Show (B, "B");
33
34
       Put_Line ("-----");
35
36
       -- When A and B get out of scope::
37
38
       - -
          Finalize (A);
       - -
39
           Finalize (B);
       - -
40
41
   end Show Controlled Types;
42
```

Code block metadata

Runtime output

```
Initialize: ID => 1
Initialize: ID => 2
```

```
Obj. A: ID => 1
Obj. B: ID => 2
- - - - - - - - -
A := B;
           ID => 1
Finalize:
Adjust:
           ID => 3
    (Previous ID => 2)
Obj. A: ID => 3
Obj. B: ID =>
              2
           ID => 2
Finalize:
          ID => 3
Finalize:
```

In contrast to the previous versions of the Simple_Controlled_Types package, type T now has an Id component. Moreover, we use a protected object Id_Gen that provides us with a unique ID to keep track of each controlled object. Basically, we assign an ID to each controlled object (right after it is created) via the call to Initialize. Similarly, this ID is updated via the calls to Adjust. Besides, we now have a Show procedure that displays the ID of a controlled object.

When running the application, we see that the calls to Initialize, Adjust and Finalize happen as expected. In addition, we see the objects' ID, which we will now analyze in order to understand how each object is actually handled.

First, we see the two calls to Initialize for objects A and B. Object A's ID is 1, and object B's ID is 2. This is later confirmed by the calls to Show.

The A := B assignment triggers two procedure calls: a call to Finalize (A) and a call to Adjust (A). In fact, this assignment can be described as follows:

- 1. Finalize (A) is called before the actual copy;
- 2. B's data is copied to object A;
- 3. Adjust (A) is called after that copy.

We can confirm this via the object ID: the object we handle in the call to Finalize (A) has an ID of 1, and the object we handle in the call to Adjust (A) has an ID of 2 (which originates from the copy of B to A) and is later changed (*adjusted*) to 3. Again, we can verify the correct IDs by looking at the output of the calls to Show.

Note that the call to Finalize (A) (before the copy of B's data) indicates that the previous version of object A is being finalized, i.e. it's as though the original object A is going to be destroyed and its contents are going to be lost. Actually, the object's contents are just overwritten, but the call to Finalize allows us to make proper adjustments to the object before the previous information is lost.

Finally, the new version of object A (the one whose ID is 3) and object B are finalized via the calls to Finalize (A) and Finalize (B) before the Show_Controlled_Types procedure ends.

18.2 Initialization

In this section, we cover some details about the initialization of controlled types. Most of those details are related to the initialization order. In principle, as stated in the Ada Reference Manual, "Initialize and other initialization operations are done in an arbitrary order," except in the situations that we describe later on.

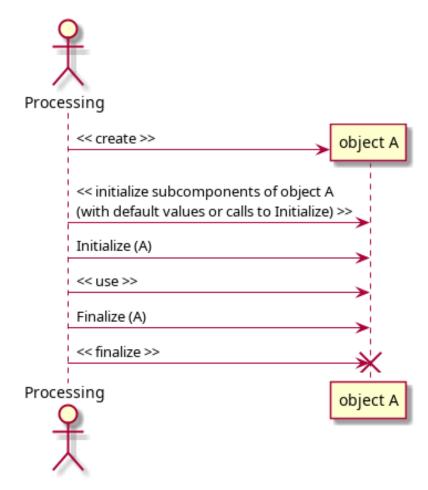
- Relevant topics
 - Assignment and Finalization³³⁷

18.2.1 Subcomponents

We've seen before that default initialization is a way of controlling the initialization of arbitrary types. In the case of controlled types, the default initialization of its subcomponents always takes places before the call to Initialize.

Similarly, a controlled type might have subcomponents of controlled types. These subcomponents are initialized by a call to the Initialize procedure of each of those controlled types.

We can visualize the lifetime as follows:



In order to see this effect, let's start by implementing two controlled types: Sub_1 and Sub_2:

Listing 11: subs.ads

```
with Ada.Finalization;
package Subs is
```

(continues on next page)

³³⁷ http://www.ada-auth.org/standards/22rm/html/RM-7-6.html

```
type Sub_1 is tagged private;
5
6
       type Sub_2 is tagged private;
7
8
   private
9
10
       type Sub_1 is new
11
         Ada.Finalization.Controlled
12
           with null record;
13
14
       overriding
15
       procedure Initialize (E : in out Sub_1);
16
17
       type Sub_2 is new
18
         Ada.Finalization.Controlled
19
           with null record;
20
21
       overriding
22
       procedure Initialize (E : in out Sub_2);
23
24
   end Subs;
25
```

Listing 12: subs.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Subs is
3
4
      procedure Initialize (E : in out Sub_1) is
5
6
      begin
          Put_Line ("Initialize: Sub_1...");
7
      end Initialize;
8
9
      procedure Initialize (E : in out Sub_2) is
10
      begin
11
          Put Line ("Initialize: Sub 2...");
12
      end Initialize;
13
14
   end Subs;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Initialization.

GControlled_Initialization

MD5: f6a7676e82294a62965157d2ffd4ae3b
```

Now, let's use those controlled types as components of a type T. In addition, let's declare an integer component I with default initialization. This is how the complete code looks like:

Listing 13: simple_controlled_types.ads

```
with Ada.Finalization;
with Subs; use Subs;
package Simple_Controlled_Types is
type T is tagged private;
procedure Dummy (E : T);
```

```
private
11
12
       function Default_Init return Integer;
13
14
       type T is new
15
         Ada.Finalization.Controlled with
16
       record
17
          S1 : Sub_1;
18
          S2 : Sub_2;
19
          I : Integer := Default_Init;
20
       end record;
21
22
23
       overriding
       procedure Initialize (E : in out T);
24
25
   end Simple_Controlled_Types;
26
```

Listing 14: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Controlled_Types is
3
4
      function Default_Init return Integer is
5
      begin
6
         Put Line ("Default Init: Integer...");
7
          return 42;
8
9
      end Default Init;
10
      procedure Dummy (E : T) is
11
12
      begin
         Put_Line ("(Dummy: T...)");
13
      end Dummy;
14
15
      procedure Initialize (E : in out T) is
16
      begin
17
          Put Line ("Initialize: T...");
18
      end Initialize;
19
20
   end Simple_Controlled_Types;
21
```

Listing 15: show_controlled_types.adb

```
with Simple_Controlled_Types;
use Simple_Controlled_Types;
procedure Show_Controlled_Types is
    A : T;
begin
    Dummy (A);
end Show_Controlled_Types;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Initialization.

→Controlled_Initialization

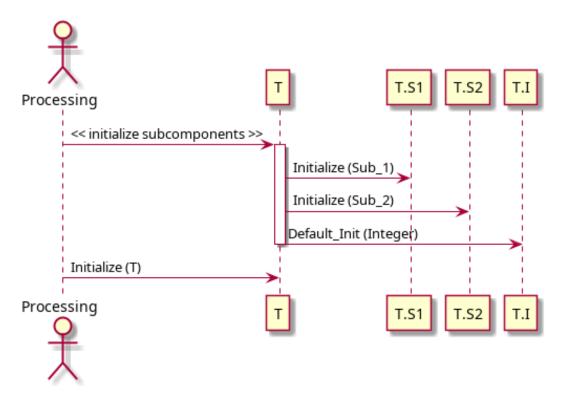
MD5: 39d0efa76c056ac8190573c86f17c890
```

Runtime output

```
Initialize: Sub_1...
Initialize: Sub_2...
Default_Init: Integer...
Initialize: T...
(Dummy: T...)
```

When we run this application, we see that the Sub_1 and Sub_2 components are initialized by calls to their respective Initialize procedures, and the I component is initialized with its default value (via a call to the Default_Init function). Finally, after all subcomponents of type T have been initialized, the Initialize procedure is called for the type T itself.

This diagram shows the initialization sequence:



18.2.2 Components with access discriminants

Record types with access discriminants are a special case. In fact, according to the Ada Reference Manual, "if an object has a component with an access discriminant constrained by a *per-object expression* (page 240), Initialize is applied to this component after any components that do not have such discriminants. For an object with several components with such a discriminant, Initialize is applied to them in order of their component declarations."

Let's see a code example. First, we implement another package with controlled types:

```
Listing 16: selections.ads
```

```
with Ada.Finalization;
package Selections is
type Selection is private;
type Selection_1 (S : access Selection) is
tagged private;
```

9

(continued from previous page)

```
type Selection_2 (S : access Selection) is
10
         tagged private;
11
12
   private
13
14
      type Selection is null record;
15
16
       type Selection_1 (S : access Selection) is new
17
         Ada.Finalization.Controlled
18
           with null record;
19
20
       overriding
21
22
      procedure Initialize
         (E : in out Selection_1);
23
24
       type Selection_2 (S : access Selection) is new
25
         Ada.Finalization.Controlled
26
           with null record;
27
28
      overriding
29
       procedure Initialize
30
         (E : in out Selection_2);
31
32
   end Selections;
33
```

Listing 17: selections.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Selections is
3
4
      procedure Initialize
5
         (E : in out Selection_1) is
6
      begin
7
          Put Line ("Initialize: Selection 1...");
8
      end Initialize;
9
10
      procedure Initialize
11
         (E : in out Selection_2) is
12
      begin
13
         Put_Line ("Initialize: Selection_2...");
14
      end Initialize;
15
16
   end Selections;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Initialization.

Gontrolled_Initialization

MD5: 01c3639ebd52d37856e77ccfeb057d1b
```

In this example, we see the declaration of the Selection_1 and Selection_2 types, which are controlled types with an access discriminant of Selection type. Now, let's use these types in the declaration of the T type from the *previous example* (page 848) and add two new components (Sel_1 and Sel_2):

Listing 18: simple_controlled_types.ads

with Ada.Finalization;

```
2
   with Subs;
                  use Subs;
3
   with Selections; use Selections;
4
5
   package Simple_Controlled_Types is
6
7
      type T (S1 : access Selection;
8
               S2 : access Selection) is
9
         tagged private;
10
11
      procedure Dummy (E : T);
12
13
   private
14
15
      function Default_Init return Integer;
16
17
       type T (S1 : access Selection;
18
               S2 : access Selection) is new
19
         Ada.Finalization.Controlled with
20
       record
21
          Sel 1 : Selection 1 (S1);
22
          Sel_2 : Selection_2 (S2);
23
          S_1 : Sub_1;
24
          Ι
               : Integer := Default_Init;
25
      end record;
26
27
28
      overriding
      procedure Initialize (E : in out T);
29
30
   end Simple_Controlled_Types;
31
```

Listing 19: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple Controlled Types is
3
4
      function Default Init return Integer is
5
      begin
6
          Put_Line ("Default_Init: Integer...");
7
          return 42;
8
      end Default_Init;
9
10
      procedure Dummy (E : T) is
11
      begin
12
         Put Line ("(Dummy: T...)");
13
      end Dummy;
14
15
      procedure Initialize (E : in out T) is
16
      begin
17
          Put Line ("Initialize: T...");
18
      end Initialize;
19
20
   end Simple_Controlled_Types;
21
```

Listing 20: show_controlled_types.adb

```
with Simple_Controlled_Types;
use Simple_Controlled_Types;
```

```
with Selections;
4
   use Selections;
5
6
   procedure Show_Controlled_Types is
7
      S1, S2 : aliased Selection;
8
      A : T (S1'Access, S2'Access);
9
   beain
10
      Dummy (A);
11
  end Show_Controlled_Types;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Initialization.

→Controlled_Initialization

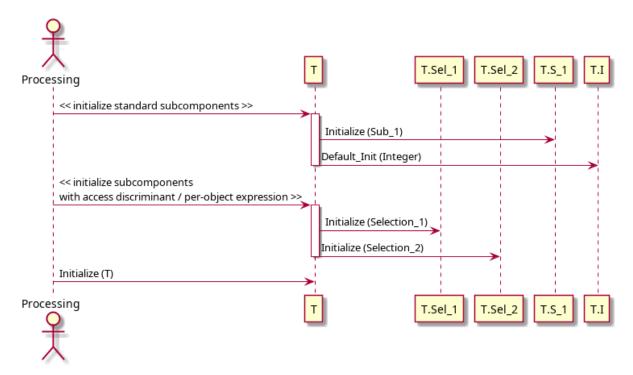
MD5: 74f507b912ab746b70aec451a9bc8f74
```

Runtime output

```
Initialize: Sub_1...
Default_Init: Integer...
Initialize: Selection_1...
Initialize: Selection_2...
Initialize: T...
(Dummy: T...)
```

When running this example, we see that all other subcomponents — to be more precise, those subcomponents that require initialization — are initialized before the Sub_1 and Sub_2 components are initialized via calls to their corresponding Initialize procedure. Note that, although Sub_1 and Sub_2 are the last components to be initialized, they are still initialized before the call to the Initialize procedure of type T.

This diagram shows the initialization sequence:



18.2.3 Task activation

Components of task types also require special treatment. According to the Ada Reference Manual, "for an allocator, any task activations follow all calls on Initialize."

As always, let's analyze an example that illustrates this. First, we implement another package called Workers with a simple task type:

Listing 21: workers.ads

```
1 package Workers is
2
3 task type Worker is
4 entry Start;
5 entry Stop;
6 end Worker;
7
8 end Workers;
```

Listing 22: workers.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Workers is
3
4
      task body Worker is
5
6
          function Init return Integer is
7
          begin
8
             Put Line ("Activating Worker task...");
9
             return 0;
10
          end Init;
11
12
          I : Integer := Init;
13
      begin
14
15
          accept Start do
16
            Put Line ("Worker.Start accepted...");
17
             I := I + 1;
18
          end Start;
19
20
          accept Stop do
21
            Put_Line ("Worker.Stop accepted...");
22
             I := I - 1;
23
          end Stop;
24
       end Worker;
25
26
   end Workers;
27
```

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Initialization. Gontrolled_Initialization MD5: 1d48a78f14a496c8cdadeab9d1bc9070

Let's extend the declaration of the T type from the *previous example* (page 851) and declare a new component of Worker type. Note that we have to change T to a limited controlled type because of this new component of task type. This is the updated code:

```
Listing 23: simple_controlled_types.ads
```

```
with Ada.Finalization;
1
2
   with Subs;
                     use Subs;
3
   with Selections; use Selections;
4
   with Workers; use Workers;
5
6
   package Simple Controlled Types is
7
8
      type T (S : access Selection) is
9
        tagged limited private;
10
11
      procedure Start_Work (E : T);
12
      procedure Stop_Work (E : T);
13
14
   private
15
16
      function Default_Init return Integer;
17
18
      type T (S : access Selection) is new
19
        Ada.Finalization.Limited_Controlled with
20
       record
21
         W
                : Worker;
22
         Sel_1 : Selection_1 (S);
23
         S1 : Sub_1;
24
                : Integer := Default_Init;
         Ι
25
      end record;
26
27
      overriding
28
      procedure Initialize (E : in out T);
29
30
   end Simple_Controlled_Types;
31
```

Listing 24: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Controlled_Types is
3
4
      function Default_Init return Integer is
5
      begin
6
          Put_Line ("Default_Init: Integer...");
7
          return 42;
8
      end Default_Init;
9
10
      procedure Start_Work (E : T) is
11
      begin
12
              Starting Worker task:
13
          E.W.Start;
14
15
      end Start_Work;
16
17
      procedure Stop_Work (E : T) is
18
       begin
19
              Stopping Worker task:
20
          E.W.Stop;
21
      end Stop_Work;
22
23
       procedure Initialize (E : in out T) is
24
25
      begin
```

```
26 Put_Line ("Initialize: T...");
27 end Initialize;
28
29 end Simple_Controlled_Types;
```

Listing 25: show_controlled_types.adb

```
with Simple_Controlled_Types;
1
   use Simple_Controlled_Types;
2
3
   with Selections; use Selections;
4
5
   procedure Show_Controlled_Types is
6
      type T_Access is access T;
7
8
      S : aliased Selection;
9
      A : constant T Access := new T (S'Access);
10
   begin
11
      Start_Work (A.all);
12
      Stop Work (A.all);
13
  end Show_Controlled_Types;
14
```

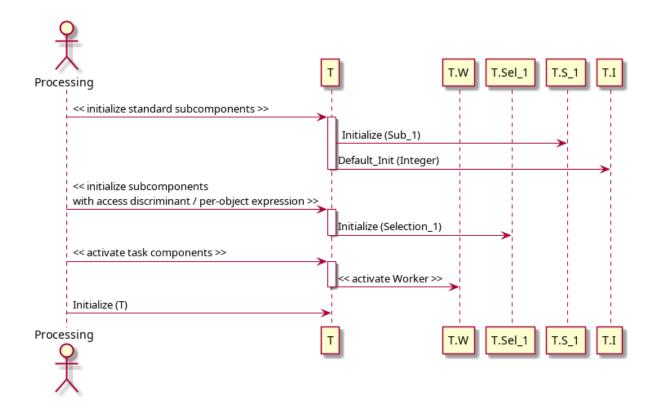
Code block metadata

Runtime output

```
Initialize: Sub_1...
Default_Init: Integer...
Initialize: Selection_1...
Initialize: T...
Activating Worker task...
Worker.Start accepted...
Worker.Stop accepted...
```

When we run this application, we see that the W component is activated only after all other subcomponents of type T have been initialized.

This diagram shows the initialization sequence:



18.3 Assignment

We already talked about *adjustments* (page 842) previously. As we already mentioned, an actual assignment is a full bit-wise copy of the entire right-hand side to the entire left-hand side, so the adjustment (via a call to Adjust) is a way to "work around" that, when necessary. In this section, we'll look into some details about the adjustment of controlled types.

Relevant topics

Assignment and Finalization³³⁸

18.3.1 Assignment using anonymous object

The Ada Reference Manual³³⁹ mentions that an anonymous object is created during the assignment of objects of controlled type. A simple A := B operation for nonlimited controlled types can be expanded to the following illustrative code:

```
procedure P is
    A, B: Some_Controlled_Type;
begin
    --
    A := B;
    --
    B_To_A_Assignment : declare
    Anon_Obj : Some_Controlled_Type;
    begin
    Anon_Obj := B;
```

(continues on next page)

³³⁸ http://www.ada-auth.org/standards/22rm/html/RM-7-6.html
 ³³⁹ http://www.ada-auth.org/standards/22rm/html/RM-7-6.html

```
Adjust (Anon_Obj);
Finalize (A);
A := Anon_Obj;
Finalize (Anon_Obj);
end B_To_A_Assignment;
end P;
```

The first assignment happens to the anonymous object Anon_Obj. After the adjustment of Anon_Obj and the finalization of the original version of A, the actual assignment to A can take place — and Anon_Obj can be discarded after it has been properly finalized. With this strategy, we have a chance to finalize the original version of A before the assignment overwrites the object.

Of course, this expanded code isn't really efficient, and the compiler has some freedom to improve the performance of the generated machine code. Whenever possible, it'll typically optimize the anonymous object out and build the object in place. (The Ada Reference Manual³⁴⁰ describes the rules when this is possible or not.)

Also, the A := Anon_0bj statement in the code above doesn't necessarily translate to an actual assignment in the generated machine code. Typically, a compiler may treat Anon_0bj as the new A and destroy the original version of A (i.e. the object that used to be A). In this case, the code becomes something like this:

```
procedure P is
  A, B: Some Controlled Type;
begin
  --A := B;
  B_To_A_Assignment : declare
     Anon Obj : Some Controlled Type;
  begin
      Anon Obj := B;
      Finalize (A);
     Adjust (Anon Obj);
      declare
         A : Some_Controlled_Type renames Anon_Obj;
      begin
         -- Now, we treat Anon_Obj as the new A.
         -- Further processing continues here...
      end:
  end B To A Assignment;
end P;
```

In some cases, the compiler is required to build the object in place. A typical example is when an object of controlled type is initialized by assigning an aggregate to it:

```
C: constant Some_Controlled_Type :=
   (Ada.Finalization.Controlled with ...);
-- C is built in place,
-- no anonymous object is used here.
```

Also, it's possible that Adjust and Finalize aren't called at all. Consider an assignment like this: A := A; In this case, since the object on both sides is the same, the compiler is allowed to simply skip the assignment and not do anything.

For more details about possible optimizations and compiler behavior, please refer to the Ada Reference $Manual^{341}$.

³⁴⁰ http://www.ada-auth.org/standards/22rm/html/RM-7-6.html

³⁴¹ http://www.ada-auth.org/standards/22rm/html/RM-7-6.html

In general, the advice is simple: use Adjust and Finalize solely for their intended purposes. In other words, don't implement extraneous side-effects into those procedures, as they might not be called at run-time.

18.3.2 Adjustment of subcomponents

In principle, the order in which components are adjusted is arbitrary. However, adjustments of subcomponents will happen before the adjustment of the component itself. The subcomponents must be adjusted before the enclosing object because the semantics of the adjustment of the whole might depend on the states of the parts (the subcomponents), so those states must already be in place.

Let's revisit a *previous code example* (page 848). First, we override the Adjust procedure of the Sub_1 and Sub_2 types from the Subs package.

Listing 26: subs.ads

```
with Ada.Finalization;
1
   package Subs is
3
4
      type Sub 1 is tagged private;
5
6
       type Sub 2 is tagged private;
7
8
   private
9
10
       type Sub 1 is new
11
         Ada.Finalization.Controlled
12
           with null record;
13
14
      overriding
15
      procedure Initialize (E : in out Sub_1);
16
17
      overriding
18
      procedure Adjust (E : in out Sub_1);
19
20
       overriding
21
      procedure Finalize (E : in out Sub_1);
22
23
       type Sub_2 is new
24
         Ada.Finalization.Controlled
25
           with null record;
26
27
      overriding
28
      procedure Initialize (E : in out Sub_2);
29
30
      overriding
31
      procedure Adjust (E : in out Sub_2);
32
33
      overriding
34
      procedure Finalize (E : in out Sub_2);
35
36
```

37 end Subs;

Listing 27: subs.adb

```
with Ada.Text_IO; use Ada.Text_IO;
package body Subs is
```

```
procedure Initialize (E : in out Sub_1) is
5
      begin
6
          Put_Line ("Initialize: Sub_1...");
7
      end Initialize;
8
9
      procedure Adjust (E : in out Sub_1) is
10
      beain
11
          Put_Line ("Adjust: Sub_1...");
12
      end Adjust;
13
14
      procedure Finalize (E : in out Sub_1) is
15
      begin
16
          Put_Line ("Finalize: Sub_1...");
17
18
       end Finalize;
19
      procedure Initialize (E : in out Sub_2) is
20
      begin
21
          Put_Line ("Initialize: Sub_2...");
22
      end Initialize;
23
24
      procedure Adjust (E : in out Sub_2) is
25
      begin
26
          Put_Line ("Adjust: Sub_2...");
27
       end Adjust;
28
29
      procedure Finalize (E : in out Sub_2) is
30
31
      begin
          Put_Line ("Finalize: Sub_2...");
32
      end Finalize;
33
34
   end Subs;
35
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Adjustment.

GControlled_Initialization

MD5: 110d88543a7a897ba433c90f6c2a881c
```

Next, we override the Adjust procedure of the T type from the Simple_Controlled_Types package:

Listing 28: simple_controlled_types.ads

```
with Ada.Finalization;
1
2
   with Subs; use Subs;
3
4
   package Simple Controlled Types is
5
6
      type T is tagged private;
7
8
      procedure Dummy (E : T);
9
10
   private
11
12
       function Default_Init return Integer;
13
14
       type T is new
15
         Ada.Finalization.Controlled with
16
       record
17
          S1 : Sub 1;
18
```

```
S2 : Sub_2;
19
          I : Integer := Default_Init;
20
       end record;
21
22
      overriding
23
      procedure Initialize (E : in out T);
24
25
      overriding
26
      procedure Adjust (E : in out T);
27
28
      overriding
29
       procedure Finalize (E : in out T);
30
31
   end Simple_Controlled_Types;
32
```

Listing 29: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Controlled_Types is
3
4
      function Default_Init return Integer is
5
      begin
6
          Put_Line ("Default_Init: Integer...");
7
          return 42;
8
      end Default Init;
9
10
      procedure Dummy (E : T) is
11
12
      begin
          Put_Line ("(Dummy: T...)");
13
14
      end Dummy;
15
      procedure Initialize (E : in out T) is
16
      begin
17
          Put Line ("Initialize: T...");
18
      end Initialize;
19
20
      procedure Adjust (E : in out T) is
21
      begin
22
          Put_Line ("Adjust: T...");
23
      end Adjust;
24
25
      procedure Finalize (E : in out T) is
26
27
      begin
          Put_Line ("Finalize: T...");
28
      end Finalize;
29
30
   end Simple Controlled Types;
31
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Adjustment.

→Controlled_Initialization

MD5: 9fb392305df70734994cffe612cb3869
```

Finally, this is the main application:

Listing 30: show controlled types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
```

```
with Simple_Controlled_Types;
3
   use Simple_Controlled_Types;
4
5
   procedure Show_Controlled_Types is
6
      A, B : T;
7
   begin
8
      Dummy (A);
9
10
      Put_Line ("-----");
11
      Put_Line ("A := B");
12
      A := B;
13
      Put_Line ("-----");
14
   end Show_Controlled_Types;
15
```

Code block metadata

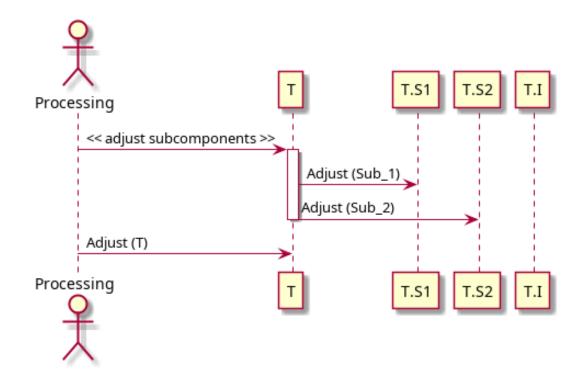
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Adjustment. →Controlled_Initialization MD5: 1ceaa50cbb18b9f1f997246a614e3a90

Runtime output

Initialize: Sub_1... Initialize: Sub 2... Default Init: Integer... Initialize: T... Initialize: Sub 1... Initialize: Sub 2... Default Init: Integer... Initialize: T... (Dummy: T...) - - - - - - - - - -A := B Finalize: T... Finalize: Sub_2... Finalize: Sub_1... Adjust: Sub 1... Adjust: Sub_2... Adjust: T... - - - - - - - - - -Finalize: T... Finalize: Sub 2... Finalize: Sub_1... Finalize: T... Finalize: Sub 2... Finalize: Sub 1...

When running this code, we see that the S1 and S2 components are adjusted before the adjustment of the parent type T takes place.

This diagram shows the adjustment sequence:



18.4 Finalization

We mentioned finalization — and the Finalize procedure — at the *beginning of the chapter* (page 838). In this section, we discuss the topic in more detail.

Relevant topics

- Assignment and Finalization³⁴²
- Completion and Finalization³⁴³

18.4.1 Normal and abnormal completion

When a subprogram has just executed its last statement, normal completion of this subprogram has been reached. At this point, finalization starts. In the case of controlled objects, this means that the Finalize procedure is called for those objects. (As we've already seen *an example of normal completion* (page 840) at the beginning of the chapter, we won't repeat it here, as we assume you are already familiar with the concept.)

When an exception is raised or due to an abort, however, a subprogram has an abnormal completion. We discuss more about exception handling and finalization *later on* (page 873).

18.4.2 Finalization via unchecked deallocation

When performing unchecked deallocation of a controlled type, the Finalize procedure is called right before the actual memory for the controlled object is deallocated.

Let's see a simple example:

³⁴² http://www.ada-auth.org/standards/22rm/html/RM-7-6.html

³⁴³ http://www.ada-auth.org/standards/22rm/html/RM-7-6-1.html

Listing 31: simple_controlled_types.ads

```
with Ada.Finalization;
1
   with Ada.Unchecked_Deallocation;
2
3
   package Simple_Controlled_Types is
4
5
      type T is tagged private;
6
7
      procedure Dummy (E : T);
8
9
      type T_Access is access T;
10
11
      procedure Free (A : in out T_Access);
12
13
   private
14
15
      type T is new
16
        Ada.Finalization.Controlled
17
           with null record;
18
19
      overriding
20
      procedure Finalize (E : in out T);
21
22
      procedure Free_T_Access is
23
         new Ada.Unchecked_Deallocation
24
           (Object => T,
25
            Name
                  => T_Access);
26
27
      procedure Free (A : in out T_Access)
28
         renames Free_T_Access;
29
30
   end Simple_Controlled_Types;
31
```

Listing 32: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple_Controlled_Types is
3
4
      procedure Dummy (E : T) is
5
      begin
6
         Put_Line ("(Dummy T...)");
7
      end Dummy;
8
9
      procedure Finalize (E : in out T) is
10
11
      begin
         Put_Line ("Finalize T...");
12
      end Finalize;
13
14
   end Simple_Controlled_Types;
15
```

Listing 33: show_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
with Simple_Controlled_Types;
use Simple_Controlled_Types;
procedure Show_Controlled_Types is
A : T_Access := new T;
```

```
begin
8
      Dummy (A.all);
9
10
      Free (A);
11
       -- At this point, Finalize (A.all)
12
       -- will be called before the actual
13
       -- deallocation.
14
15
      Put_Line ("We've just freed A.");
16
   end Show_Controlled_Types;
17
```

Code block metadata

Runtime output

```
(Dummy T...)
Finalize T...
We've just freed A.
```

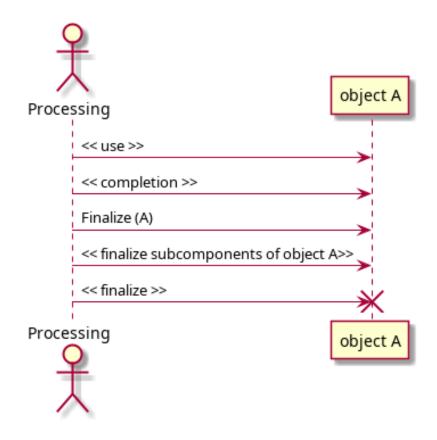
In this example, we see that a call to Finalize (for type T) is triggered by the call to Free for the A object — at this point, we haven't reached the end of the main procedure (Show_Controlled_Types) yet. After the call to Free, the object originally referenced by A has been completely finalized — and deallocated.

When the main procedure completes (after the call to Put_Line in that procedure), we would normally see the calls to Finalize for controlled objects. However, at this point, we obviously don't have a second call to the Finalize procedure for type T, as the object referenced by A has already been finalized and freed.

18.4.3 Subcomponents

As we've seen in the section about *initialization of subcomponents* (page 848), subcomponents of a controlled type are initialized by a call to their corresponding Initialize procedure before the call to Initialize for the parent controlled type. In the case of finalization, the reverse order is applied: first, finalization of the parent type takes place, and then the finalization of the subcomponents.

We can visualize the lifetime as follows:



Let's show a code example by revisiting the previous implementation of the controlled types Sub_1 and Sub_2, and adapting it:

```
Listing 34: subs.ads
```

```
with Ada.Finalization;
1
2
   package Subs is
3
4
       type Sub_1 is tagged private;
5
6
       type Sub 2 is tagged private;
7
8
   private
9
10
       type Sub_1 is new
11
         Ada.Finalization.Controlled
12
           with null record;
13
14
       overriding
15
       procedure Finalize (E : in out Sub_1);
16
17
       type Sub_2 is new
18
         Ada.Finalization.Controlled
19
           with null record;
20
21
       overriding
22
       procedure Finalize (E : in out Sub_2);
23
24
   end Subs;
25
```

Listing 35: subs.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Subs is
3
4
      procedure Finalize (E : in out Sub_1) is
5
      begin
6
          Put_Line ("Finalize: Sub_1...");
7
      end Finalize;
8
9
      procedure Finalize (E : in out Sub_2) is
10
      begin
11
         Put_Line ("Finalize: Sub_2...");
12
      end Finalize;
13
14
   end Subs;
15
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Finalization.

→Controlled_Initialization

MD5: 565f0b13586c08e0cdfdc119bcb28780
```

Now, let's use those controlled types as components of a type T:

```
Listing 36: simple_controlled_types.ads
```

```
with Ada.Finalization;
1
2
   with Subs; use Subs;
3
4
   package Simple_Controlled_Types is
5
6
       type T is tagged private;
7
8
       procedure Dummy (E : T);
9
10
   private
11
12
       type T is new
13
         Ada.Finalization.Controlled with
14
15
       record
          S1 : Sub_1;
16
17
          S2 : Sub_2;
       end record;
18
19
20
       overriding
       procedure Finalize (E : in out T);
21
22
   end Simple_Controlled_Types;
23
```

Listing 37: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
package body Simple_Controlled_Types is
procedure Dummy (E : T) is
begin
Put_Line ("(Dummy: T...)");
```

```
8 end Dummy;
9
10 procedure Finalize (E : in out T) is
11 begin
12 Put_Line ("Finalize: T...");
13 end Finalize;
14
15 end Simple_Controlled_Types;
```

Listing 38: show_controlled_types.adb

```
1 with Simple_Controlled_Types;
2 use Simple_Controlled_Types;
3 
4 procedure Show_Controlled_Types is
5 A : T;
6 begin
7 Dummy (A);
8 end Show_Controlled_Types;
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Finalization.

→Controlled_Initialization

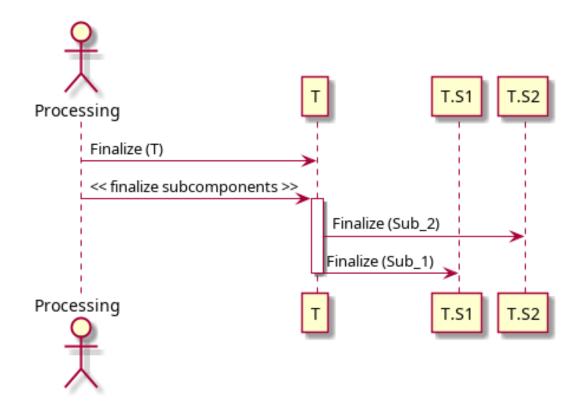
MD5: 6feecb7c544f340bf4841034d7ab5f71
```

Runtime output

(Dummy: T...) Finalize: T... Finalize: Sub_2... Finalize: Sub_1...

When we run this application, we see that the Finalize procedure is called for the type T itself — as the first step of the finalization of type T. Then, the Sub_2 and Sub_1 components are finalized by calls to their respective Finalize procedures.

This diagram shows the finalization sequence:



18.4.4 Components with access discriminants

We already discussed the *initialization of components with access discriminants constrained by a per-object expression* (page 851). In the case of the finalization of such components, they are finalized before any components that do not fall into this category — in the reverse order of their component declarations — but after the finalization of the parent type.

Let's revisit a *previous code example* (page 851) and adapt it to demonstrate the finalization of components with access discriminants. First, we implement another package with controlled types:

```
Listing 39: selections.ads
```

```
with Ada.Finalization;
1
2
   package Selections is
3
4
      type Selection is private;
5
6
      type Selection_1 (S : access Selection) is
7
         tagged private;
8
9
       type Selection_2 (S : access Selection) is
10
         tagged private;
11
12
   private
13
14
      type Selection is null record;
15
16
       type Selection_1 (S : access Selection) is new
17
         Ada.Finalization.Controlled
18
           with null record;
19
20
      overriding
21
      procedure Finalize
22
```

```
(E : in out Selection_1);
23
24
       type Selection_2 (S : access Selection) is new
25
         Ada.Finalization.Controlled
26
           with null record;
27
28
      overriding
29
      procedure Finalize
30
         (E : in out Selection_2);
31
32
   end Selections;
33
```

Listing 40: selections.adb

```
with Ada.Text IO; use Ada.Text IO;
1
2
   package body Selections is
3
4
      procedure Finalize
5
        (E : in out Selection 1) is
6
      begin
7
          Put_Line ("Finalize: Selection_1...");
8
      end Finalize;
9
10
      procedure Finalize
11
        (E : in out Selection 2) is
12
      begin
13
          Put Line ("Finalize: Selection 2...");
14
      end Finalize;
15
16
   end Selections;
17
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Finalization.

←Controlled_Initialization

MD5: d1d35eb7ea62742fb130fbf05d898989
```

In this example, we see the declaration of the Selection_1 and Selection_2 types, which are controlled types with an access discriminant of Selection type. Now, let's use these types in the declaration of a type T and add two new components — Sel_1 and Sel_2:

Listing 41: simple_controlled_types.ads

```
with Ada.Finalization;
1
2
   with Subs;
                      use Subs;
3
   with Selections; use Selections;
4
5
   package Simple_Controlled_Types is
6
7
       type T (S1 : access Selection;
8
9
               S2 : access Selection) is
         tagged private;
10
11
      procedure Dummy (E : T);
12
13
   private
14
15
       type T (S1 : access Selection;
16
```

```
S2 : access Selection) is new
17
         Ada.Finalization.Controlled with
18
       record
19
          Sel_1 : Selection_1 (S1);
20
          Sel_2 : Selection_2 (S2);
21
          S_1 : Sub_1;
22
      end record;
23
24
      overriding
25
      procedure Finalize (E : in out T);
26
27
   end Simple_Controlled_Types;
28
```

Listing 42: simple_controlled_types.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body Simple Controlled Types is
3
4
      procedure Dummy (E : T) is
5
      begin
6
         Put_Line ("(Dummy: T...)");
7
      end Dummy;
8
9
      procedure Finalize (E : in out T) is
10
      begin
11
         Put Line ("Finalize: T...");
12
      end Finalize;
13
14
   end Simple_Controlled_Types;
15
```

Listing 43: show_controlled_types.adb

```
with Simple_Controlled_Types;
1
2
   use Simple_Controlled_Types;
3
   with Selections;
4
   use Selections;
5
6
   procedure Show_Controlled_Types is
7
      S1, S2 : aliased Selection;
8
      A : T (S1'Access, S2'Access);
9
  begin
10
11
      Dummy (A);
  end Show_Controlled_Types;
12
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Finalization.

→Controlled_Initialization

MD5: e421a750f11ade3b4df98569c71b904a
```

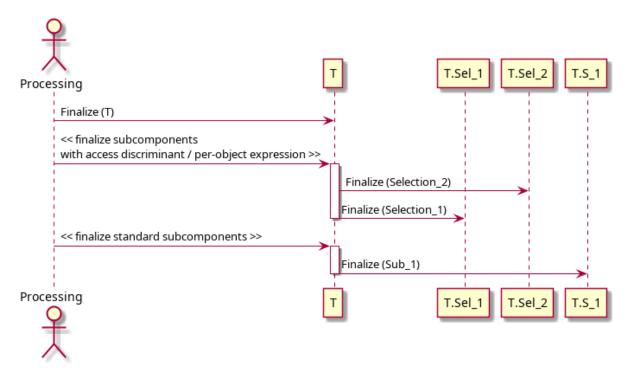
Runtime output

```
(Dummy: T...)
Finalize: T...
Finalize: Selection_2...
Finalize: Selection_1...
Finalize: Sub_1...
```

When we run this example, we see that the Finalize procedure of type T is called as

the first step. Then, the Finalize procedure is called for the components with an access discriminant constrained by a *per-object expression* (page 240) — in this case, Sel_2 and Sel_1 (of Selection_2 and Selection_1 types, respectively). Finally, the Sub_1 component is finalized.

This diagram shows the finalization sequence:



18.5 Controlled Types and Exception Handling

In the previous section, we mainly focused on the normal completion of controlled types. However, when control is transferred out of the normal execution path due to an abort or an exception being raised, we speak of abnormal completion. In this section, we focus on those cases.

Let's start with a simple example:

```
Listing 44: simple_controlled_types.ads
```

```
with Ada.Finalization;
1
2
   package Simple Controlled Types is
3
4
      type T is tagged private;
5
6
      procedure Dummy (E : T);
7
8
   private
9
10
       type T is new
11
         Ada.Finalization.Controlled
12
           with null record;
13
14
      overriding
15
      procedure Initialize (E : in out T);
16
17
      overriding
18
```

```
procedure Finalize (E : in out T);
end Simple_Controlled_Types;
```

Listing 45: simple_controlled_types.adb

```
with Ada.Text IO; use Ada.Text IO;
1
   package body Simple_Controlled_Types is
3
4
      procedure Dummy (E : T) is
5
      begin
6
         Put_Line ("(Dummy...)");
7
      end Dummy;
8
9
      procedure Initialize (E : in out T) is
10
      begin
11
         Put Line ("Initialize...");
12
      end Initialize;
13
14
      procedure Finalize (E : in out T) is
15
      begin
16
         Put Line ("Finalize...");
17
      end Finalize;
18
19
   end Simple_Controlled_Types;
20
```

Listing 46: show_simple_exception.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Simple Controlled Types;
3
   use Simple_Controlled_Types;
4
5
   procedure Show_Simple_Exception is
6
       A : T;
7
8
       function Int_Last return Integer is
9
         (Integer'Last);
10
11
       Cnt : Positive := Int_Last;
12
   begin
13
       Cnt := Cnt + 1;
14
15
       Dummy (A);
16
17
       Put_Line (Cnt'Image);
18
19
       -- When A is about to get out of
20
       - -
          scope:
21
22
       - -
          Finalize (A);
       - -
23
24
   end Show_Simple_Exception;
25
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Exception_

→Handling.Simple_Exception

MD5: 9461f420f091f058e6ea1ee419b2a5c6
```

Runtime output

```
Initialize...
Finalize...
```

raised CONSTRAINT_ERROR : show_simple_exception.adb:14 overflow check failed

In this example, we're forcing an overflow to happen in the Show_Simple_Exception by adding one to the integer variable Cnt, which already has the value **Integer**'Last. The corresponding *overflow check* (page 519) raises the Constraint_Error.

However, *before* this exception is raised, the finalization of the controlled object A is performed. In this sense, we have normal completion of the controlled type — even though an exception is being raised.

For further reading...

We already talked about the *allocation check* (page 523), which may raise a Program_Error exception. In the code example for that section, we used controlled types. Feel free to revisit the example.

Relevant topics

Completion and Finalization³⁴⁴

18.5.1 Exception raising in Initialize

If an exception is raised in the Initialize procedure, we have abnormal completion. Let's see an example:

Listing 47:	ct	initialize	exception.ads

```
with Ada.Finalization:
1
2
   package CT Initialize Exception is
3
4
      type T is tagged private;
5
6
      procedure Dummy (E : T);
7
8
   private
9
10
      type T is new
11
         Ada.Finalization.Controlled
12
           with null record;
13
14
       overriding
15
      procedure Initialize (E : in out T);
16
17
       overriding
18
      procedure Finalize (E : in out T);
19
20
   end CT Initialize Exception;
21
```

³⁴⁴ http://www.ada-auth.org/standards/22rm/html/RM-7-6-1.html

Listing 48: ct_initialize_exception.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   package body CT Initialize Exception is
3
4
       function Int Last return Integer is
5
         (Integer'Last);
6
7
      Cnt : Positive := Int Last;
8
9
      procedure Dummy (E : T) is
10
      begin
11
          Put_Line ("(Dummy...)");
12
      end Dummy;
13
14
      procedure Initialize (E : in out T) is
15
       begin
16
          Put_Line ("Initialize...");
17
          Cnt := Cnt + 1;
18
      end Initialize;
19
20
      procedure Finalize (E : in out T) is
21
      begin
22
          Put_Line ("Finalize...");
23
      end Finalize;
24
25
   end CT Initialize Exception;
26
```

Listing 49: show_initialize_exception.adb

```
with CT_Initialize_Exception;
use CT_Initialize_Exception;

procedure Show_Initialize_Exception is
    A : T;
    begin
    Dummy (A);
    end Show_Initialize_Exception;
```

Code block metadata

Runtime output

```
Initialize...
```

raised CONSTRAINT_ERROR : ct_initialize_exception.adb:18 overflow check failed

In the Show_Initialize_Exception procedure, we declare an object A of controlled type T. As we know, this declaration triggers a call to the Initialize procedure that we've implemented in the body of the CT_Initialize_Exception package. In the Initialize procedure, we're forcing an overflow to happen — by adding one to the Cnt variable, which already has the **Integer**'Last value.

This is an example of abnormal completion, as the control is transferred out of the Initialize procedure, and the corresponding Finalize procedure is never called for object A.

18.5.2 Bounded errors of controlled types

Bounded errors (page 506) are an important topic when talking about exception and controlled types. In general, if an exception is raised in the Adjust or Finalize procedure, this is considered a bounded error. If the bounded error is detected, the Program_Error exception is raised.

Note that the original exception raised in the Adjust or Finalize procedures could be any possible exception. For example, one of those procedures could raise a Constraint_Error exception. However, the actual exception that is raised at runtime is the Program_Error exception. This is because the bounded error, which raises the Program_Error exception, is more severe than the original exception coming from those procedures.

(The behavior is different when the Adjust or Finalize procedure is called explicitly, as we'll see later.)

Not every exception raised during an operation on controlled types is considered a bounded error. In fact, the case we've seen before, an *exception raised in the Initialize procedure* (page 875) is not a bounded error.

Here's a code example of a Constraint_Error exception being raised in the Finalize procedure:

Listing !	50:	ct	finalize	exception.ads

```
with Ada.Finalization;
1
2
   package CT Finalize Exception is
3
4
      type T is tagged private;
5
6
      procedure Dummy (E : T);
7
8
      procedure Reset Counter;
9
10
   private
11
12
       type T is new
13
         Ada.Finalization.Controlled
14
           with null record;
15
16
       overriding
17
      procedure Initialize (E : in out T);
18
19
      overriding
20
      procedure Adjust (E : in out T);
21
22
      overriding
23
      procedure Finalize (E : in out T);
24
25
   end CT_Finalize_Exception;
26
```

Listing 51: ct_finalize_exception.adb

```
with Ada.Text_I0; use Ada.Text_I0;
package body CT_Finalize_Exception is
Cnt : Integer := Integer'Last;
procedure Dummy (E : T) is
begin
Put_Line ("(Dummy...)");
```

```
end Dummy;
10
11
       procedure Initialize (E : in out T) is
12
13
       begin
          Put_Line ("Initialize...");
14
       end Initialize;
15
16
       overriding
17
       procedure Adjust (E : in out T) is
18
       begin
19
          Put_Line ("Adjust...");
20
       end Adjust;
21
22
       procedure Finalize (E : in out T) is
23
24
       begin
          Put_Line ("Finalize...");
25
          Cnt := Cnt + 1;
26
       end Finalize;
27
28
       procedure Reset_Counter is
29
       begin
30
          Cnt := 0;
31
       end Reset_Counter;
32
33
   end CT_Finalize_Exception;
34
```

Listing 52: show_finalize_exception.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with CT_Finalize_Exception;
3
   use CT_Finalize_Exception;
4
5
   procedure Show_Finalize_Exception is
6
      A, B : T;
7
   begin
8
      Dummy (A);
9
10
       -- When A is about to get out of
11
      -- scope:
12
       - -
13
      -- Finalize (A);
14
15
   end Show_Finalize_Exception;
16
```

Code block metadata

Runtime output

In this example, we're again forcing an overflow to happen (by adding one to the integer variable Cnt), this time in the Finalize procedure. When this procedure is implicitly called — when object A is about to get out of scope in the Show_Finalize_Exception procedure — the Constraint_Error exception is raised.

As we've just seen, having an exception be raised during an implicit call to the Finalize procedure is a bounded error. Therefore, we see that the Program_Error exception is raised at runtime instead of the original Constraint_Error exception.

As we hinted in the beginning, when the Adjust or the Finalize procedure is called *explic-itly*, the exception raised in that procedure is *not* considered a bounded error. In this case, the original exception is raised.

To show an example of such an explicit call, let's first move the overriden procedures for type T (Initialize, Adjust and Finalize) out of the private part of the package CT_Finalize_Exception, so they are now visible to clients. This allows us to call the Finalize procedure explicitly:

Listing 53: ct_finalize_exception.ads

```
with Ada.Finalization:
1
2
   package CT_Finalize_Exception is
3
4
       type T is new
5
         Ada.Finalization.Controlled
6
      with null record;
7
8
9
       overriding
       procedure Initialize (E : in out T);
10
11
      overriding
12
      procedure Adjust (E : in out T);
13
14
       overriding
15
      procedure Finalize (E : in out T);
16
17
      procedure Dummy (E : T);
18
19
      procedure Reset Counter;
20
21
   end CT Finalize Exception;
22
```

Listing 54: show_finalize_exception.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with CT_Finalize_Exception;
3
   use CT_Finalize_Exception;
4
5
   procedure Show_Finalize_Exception is
6
      A : T;
7
   begin
8
      Dummy (A);
9
10
      Finalize (A);
11
12
      Put_Line ("After Finalize");
13
   exception
14
      when Constraint_Error =>
15
          Put Line
16
            ("Constraint Error is being handled...");
17
```

18 Reset_Counter; 19 end Show_Finalize_Exception;

Code block metadata

Runtime output

```
Initialize...
(Dummy...)
Finalize...
Constraint_Error is being handled...
Finalize...
```

Now, we're calling the Finalize procedure explicitly in the Show_Finalize_Exception procedure. As we know, due to the operation on I in the Finalize procedure, the Constraint_Error exception is raised in the procedure. Because we're handling this exception in the Show_Finalize_Exception procedure, we see the corresponding user message ("Constraint_Error is being handled...") at runtime.

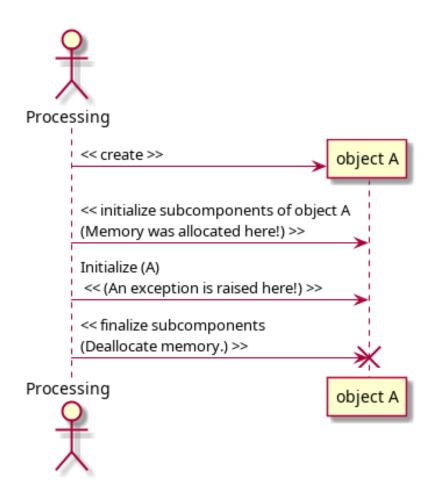
(Note that in the exception handling block, we're calling the Reset_Counter procedure. This prevents Constraint_Error from being raised in the next call to Finalize.)

18.5.3 Memory allocation and exceptions

When a memory block is allocated for controlled types and a bounded error occurs, there is no guarantee that this memory block will be deallocated. Roughly speaking, the compiler has the freedom — but not the obligation — to generate appropriate calls to Finalize, which may deallocate memory blocks.

For example, we've seen that *subcomponents of controlled type* (page 848) of a controlled object A are initialized before the initialization of object A takes place. Because memory might have been allocated for the subcomponents, the compiler can insert code that attempts to finalize those subcomponents, which in turn deallocates the memory blocks (if they were allocated in the first place).

We can visualize this strategy in the following diagram:



This strategy (of finalizing subcomponents that haven't raised exceptions) prevents memory leaks. However, this behavior very much depends on the compiler implementation. The Ada Reference Manual³⁴⁵ delineates (in the "Implementation Permissions" section) the cases where the compiler is allowed — but not required — to finalize objects when exceptions are raised.

Because the actual behavior isn't defined, custom implementation of Adjust and Finalize procedures for controlled types should be designed very carefully in order to avoid exceptions, especially when memory is allocated in the Initialize procedure.

18.6 Applications of Controlled Types

In this section, we discuss applications of controlled types. In this context, it's important to remember that controlled types have an associated overhead, which can become nonnegligible depending in which context the controlled objects are used. However, there are applications where utilizing controlled types is the best approach.

(Note that this overhead we've just mentioned is not specific to Ada. In fact, types similar to controlled types will be relatively expensive in any programming language. As an example, destructors in C++ may require a similar maintenance of state at run-time.)

³⁴⁵ http://www.ada-auth.org/standards/22rm/html/RM-7-6-1.html

18.6.1 Encapsulating access type handling

Previously, when discussing design strategies for access types (page 669), we saw an example on using *limited controlled types to encapsulate access types* (page 672).

A more generalized example is the one of an unbounded stack. Because it's unbounded, it allows for increasing the stack's size *on demand*. We can implement this kind of stack by using access types. Let's look at a simple (unoptimized) implementation:

Listing 55: unbounded stacks.ads

```
with Ada.Finalization:
1
2
   generic
3
      Default Chunk Size : Positive := 5;
4
       type Element is private;
5
   package Unbounded_Stacks is
6
      Stack Underflow : exception;
8
9
       type Unbounded Stack is private;
10
11
       procedure Push (S : in out Unbounded Stack;
12
                        Ε:
                                    Element);
13
14
       function Pop (S : in out Unbounded Stack)
15
                      return Element;
16
17
       function Is_Empty (S : Unbounded_Stack)
18
                           return Boolean;
19
20
   private
21
22
       type Element Array is
23
         array (Positive range <>) of
24
           Element;
25
26
       type Element_Array_Access is
27
         access Element_Array;
28
29
       type Unbounded Stack is new
30
         Ada.Finalization.Controlled with
31
          record
32
             Chunk Size : Positive
33
               := Default Chunk Size;
34
                         : Element_Array_Access;
             Data
35
             Тор
                         : Natural := 0;
36
          end record;
37
38
      procedure Initialize
39
         (S : in out Unbounded_Stack);
40
41
      procedure Adjust
42
         (S : in out Unbounded_Stack);
43
44
      procedure Finalize
45
         (S : in out Unbounded Stack);
46
47
   end Unbounded_Stacks;
48
```

Listing 56: unbounded_stacks.adb

```
with Ada.Text_IO; use Ada.Text_IO;
1
2
   with Ada.Unchecked_Deallocation;
3
4
   package body Unbounded Stacks is
5
6
7
          LOCAL SUBPROGRAMS
       - -
8
9
10
      procedure Free is
11
         new Ada.Unchecked_Deallocation
12
           (Object => Element_Array,
13
                  => Element_Array_Access);
            Name
14
15
      function Is_Full (S : Unbounded_Stack)
16
                          return Boolean is
17
      begin
18
          return S.Top = S.Data'Last;
19
      end Is_Full;
20
21
      procedure Reallocate_Data
22
                     : in out Element_Array_Access;
         (To
23
                           Element_Array_Access;
          From
                     :
24
                              Positive;
          Max Last
                   :
25
          Valid Last :
                              Positive) is
26
      begin
27
         To := new Element_Array (1 .. Max_Last);
28
29
          for I in 1 .. Valid_Last loop
30
             To (I) := From (I);
31
          end loop;
32
      end Reallocate_Data;
33
34
      procedure Increase_Size
35
         (S : in out Unbounded_Stack)
36
       is
37
          Old_Data : Element_Array_Access := S.Data;
38
39
          Old_Last : constant Positive
40
                      := Old_Data'Last;
          New_Last : constant Positive
41
                      := Old_Data'Last + S.Chunk_Size;
42
      begin
43
          Put_Line ("Increasing Unbounded_Stack "
44
                    & "(1 .. "
45
                    & Old_Last'Image
46
                     & ") to (1 ...
47
                     & New_Last'Image
48
                     & ")");
49
50
          Reallocate_Data
51
                       => S.Data,
52
            (To
             From
                         => Old_Data,
53
             Max_Last
                        => New_Last,
54
             Valid_Last => S.Top);
55
56
          Free (Old_Data);
57
      end Increase_Size;
58
59
```

60

(continued from previous page)

```
- -
           SUBPROGRAMS
61
62
        - -
63
       procedure Push (S : in out Unbounded_Stack;
64
                         E: Element) is
65
       begin
66
           if Is_Full (S) then
67
              Increase_Size (S);
68
           end if;
69
70
           S.Top := S.Top + 1;
71
           S.Data (S.Top) := E;
72
       end Push;
73
74
       function Pop (S : in out Unbounded_Stack)
75
                       return Element is
76
       begin
77
           return E : Element do
78
              if Is_Empty (S) then
79
                  raise Stack_Underflow;
80
              end if;
81
82
              E := S.Data (S.Top);
83
              S.Top := S.Top - 1;
84
           end return;
85
       end Pop;
86
87
       function Is_Empty (S : Unbounded_Stack)
88
                             return Boolean is
89
       begin
90
           return S.Top = 0;
91
       end Is_Empty;
92
93
94
           PRIVATE SUBPROGRAMS
95
        - -
96
97
       procedure Initialize
98
          (S : in out Unbounded_Stack)
99
       is
100
           Last : constant Positive
101
                   := S.Chunk_Size;
102
       begin
103
           Put Line ("Initializing Unbounded Stack "
104
                      & "(1 .. "
105
                      & Last'Image
106
                      & ")");
107
           S.Data := new Element_Array
108
                             (1 .. S.Chunk_Size);
109
       end Initialize;
110
111
       procedure Allocate Duplicate Data
112
          (S : in out Unbounded_Stack)
113
       is
114
           Last : constant Positive
115
                   := S.Data'Last;
116
       begin
117
           Put_Line ("Duplicating data for new "
118
                      & "Unbounded Stack (1 ..
119
                      & Last'Image
120
```

```
& ")");
121
122
           Reallocate_Data
123
                           => S.Data,
             (To
124
              From
                          => S.Data,
125
              Max_Last => Last,
126
              Valid_Last => S.Top);
127
        end Allocate_Duplicate_Data;
128
129
       procedure Adjust
130
          (S : in out Unbounded_Stack)
131
        is
132
133
       begin
           Put_Line ("Adjusting Unbounded_Stack...");
134
           Allocate_Duplicate_Data (S);
135
        end Adjust;
136
137
       procedure Finalize
138
          (S : in out Unbounded_Stack)
139
        is
140
           Last : constant Positive
141
                   := S.Data'Last;
142
       begin
143
           Put_Line ("Finalizing Unbounded_Stack "
144
                      & "(1 .. "
145
                      & Last'Image
146
                      & ")");
147
           if S.Data /= null then
148
             Free (S.Data);
149
           end if;
150
       end Finalize;
151
152
    end Unbounded_Stacks;
153
```

Listing 57: show_unbounded_stack.adb

```
with Ada.Text_I0; use Ada.Text_I0;
1
2
   with Unbounded Stacks;
3
4
   procedure Show_Unbounded_Stack is
5
6
      package Unbounded_Integer_Stacks is new
7
         Unbounded_Stacks (Element => Integer);
8
      use Unbounded_Integer_Stacks;
9
10
      procedure Print_Pop_Stack
11
              : in out Unbounded Stack;
          (S
12
           Name :
                          String)
13
       is
14
          V : Integer;
15
      begin
16
          Put_Line ("STACK: " & Name);
17
          Put ("= ");
18
          while not Is_Empty (S) loop
19
             V := Pop(S);
20
             Put (V'Image & " ");
21
          end loop;
22
          New Line;
23
       end Print_Pop_Stack;
24
25
```

```
Stack : Unbounded_Stack;
26
      Stack_2 : Unbounded_Stack;
27
28
   begin
       for I in 1 .. 10 loop
29
          Push (Stack, I);
30
       end loop;
31
32
      Stack_2 := Stack;
33
34
       for I in 11 .. 20 loop
35
          Push (Stack, I);
36
       end loop;
37
38
      Print_Pop_Stack (Stack, "Stack");
39
      Print_Pop_Stack (Stack_2, "Stack_2");
40
41
   end Show_Unbounded_Stack;
42
```

Code block metadata

```
Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Applications.

→Unbounded_Stacks

MD5: 22c795f2dfd2fbdf5468b54722d7126b
```

Runtime output

```
Initializing Unbounded Stack (1 .. 5)
Initializing Unbounded Stack (1 .. 5)
Increasing Unbounded Stack (1 .. 5) to (1 ..
                                            10)
Finalizing Unbounded_Stack (1 .. 5)
Adjusting Unbounded Stack...
Duplicating data for new Unbounded_Stack (1 .. 10)
Increasing Unbounded_Stack (1 .. 10) to (1 ..
                                             15)
Increasing Unbounded_Stack (1 .. 15) to (1 ..
                                             20)
STACK: Stack
  20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1
STACK: Stack_2
= 10 9 8 7 6 5 4 3 2 1
Finalizing Unbounded Stack (1 ...
                                10)
Finalizing Unbounded_Stack (1 ..
                                20)
```

Let's first focus on the Unbounded_Stack type from the Unbounded_Stacks package. The actual stack is implemented via the array that we allocate for the Data component. The initial allocation takes place in the Initialize procedure, which is called when an object of Unbounded_Stack type is created. The corresponding deallocation of the stack happens in the Finalize procedure.

In the Push procedure, we check whether the stack is full or not before storing a new element into the stack. If the stack is full, we call the Increase_Size procedure to *increase* the size of the array. This is actually done by calling the Reallocate_Data procedure, which allocates a new array for the stack and copies the original data to the new array.

Also, when copying an unbounded stack object to another object of this type, a call to the Adjust procedure is triggered — we do this by the assignment Stack_2 := Stack in the Show_Unbounded_Stack procedure. In the Adjust procedure, we call the Allocate_Duplicate_Data procedure to allocate a new array for the stack data and copy the data from the original stack. (Internally, the Allocate_Duplicate_Data procedure calls the Reallocate_Data procedure, which we already mentioned.)

By encapsulating the access type handling in controlled types, we can ensure that the access objects are handled correctly: no incorrect pointer usage or memory leak can happen when we use this strategy.

18.6.2 Encapsulating file handling

Controlled types can be used to encapsulate file handling, so that files are automatically created and closed. A common use-case is when a new file is expected to be created or opened when we declare the controlled object, and closed when the controlled object gets out of scope.

A simple example is the one of a logger, which we can use to write to a logfile by simple calls to Put_Line:

Listing 58: loggers.ads

```
with Ada.Text_IO; use Ada.Text_IO;
1
   with Ada.Finalization;
2
з
   package Loggers is
4
5
      type Logger (<>) is
6
         limited private;
7
8
      function Init (Filename : String)
9
                       return Logger;
10
11
      procedure Put_Line (L : Logger;
12
                             S : String);
13
14
   private
15
16
       type Logger is new
17
         Ada.Finalization.Limited Controlled with
18
19
          record
             Logfile : File_Type;
20
          end record;
21
22
      procedure Finalize
23
         (L : in out Logger);
24
25
   end Loggers;
26
```

```
Listing 59: loggers.adb
```

```
package body Loggers is
1
2
3
           SUBPROGRAMS
4
5
6
      function Init (Filename : String)
7
                       return Logger is
8
      begin
9
          return L : Logger do
10
             Create (L.Logfile, Out_File, Filename);
11
          end return;
12
      end Init;
13
14
      procedure Put_Line (L : Logger;
15
                            S : String) is
16
17
      begin
          Put_Line ("Logger: Put_Line");
18
          Put_Line (L.Logfile, S);
19
      end Put_Line;
20
21
```

```
-- PRIVATE SUBPROGRAMS
23
24
       - -
25
       procedure Finalize
26
        (L : in out Logger) is
27
       beain
28
          Put_Line ("Finalizing Logger...");
29
          if Is_Open (L.Logfile) then
30
             Close (L.Logfile);
31
          end if;
32
       end Finalize;
33
34
35
   end Loggers;
```

22

Listing 60: some_processing.adb

```
with Loggers; use Loggers;
procedure Some_Processing (Log : Logger) is
begin
Put_Line (Log, "Some processing...");
end Some_Processing;
```

Listing 61: show_logger.adb

```
with Loggers;
                          use Loggers;
1
  with Some_Processing;
2
3
  procedure Show_Logger is
4
      Log : constant Logger := Init ("report.log");
5
6
  begin
      Put_Line (Log, "Some info...");
7
      Some_Processing (Log);
8
  end Show_Logger;
```

Code block metadata

```
Some info...
Some processing...
```

Runtime output

```
Logger: Put_Line
Logger: Put_Line
Finalizing Logger...
```

The Logger type from the Loggers package has two subprograms:

- Init, which creates a logger object and creates a logfile in the background, and
- Put_Line, which writes a message to the logfile.

Note that we use the (<>) in the declaration of the Logger type to ensure that clients call the Init function. This allows us to specify the location of the logfile (as the Filename parameter).

Also, we can pass the logger to other subprograms and use it there. In this example, we

pass the logger to the Some_Processing procedure and there, we the call Put_Line using the logger object.

Finally, as soon as the logger goes out of scope, the log is automatically closed via the call to Finalize.

• For further reading....

Instead of enforcing a call to Init, we could have overridden the Initialize procedure and opened the logfile there. This approach, however, would have prevented the client from specifying the location of the logfile in a simple way. Specifying the filename as a type discriminant wouldn't work because we cannot use a string as a discriminant — as we mentioned *in a previous chapter* (page 196), we cannot use indefinite subtypes as discriminants.

If we had preferred this approach, we could generate a random name for the file in the Initialize procedure and store the file itself in a temporary directory indicated by the operating system. Alternatively, we could use the access to a string as a discriminant:

Listing 62: loggers.ads

```
with Ada.Text IO; use Ada.Text IO;
1
   with Ada.Finalization;
2
3
   package Loggers is
4
5
       type Logger (Filename : access String) is
6
         limited private;
7
8
       procedure Put Line (L : Logger;
9
                             S : String);
10
11
   private
12
13
       type Logger (Filename : access String) is new
14
         Ada.Finalization.Limited Controlled with
15
          record
16
             Logfile : File Type;
17
          end record:
18
19
       procedure Initialize
20
         (L : in out Logger);
21
22
       procedure Finalize
23
         (L : in out Logger);
24
25
   end Loggers;
26
```

```
Listing 63: loggers.adb
   package body Loggers is
      -- SUBPROGRAMS
      procedure Put_Line (L : Logger;
                           S : String) is
      begin
         Put_Line ("Logger: Put_Line");
         Put_Line (L.Logfile, S);
      end Put Line;
      -- PRIVATE SUBPROGRAMS
      - -
      procedure Initialize
       (L : in out Logger) is
20
      begin
         Create (L.Logfile,
                  Out File,
                  L.Filename.all);
      end Initialize;
      procedure Finalize
        (L : in out Logger) is
      begin
         Put Line ("Finalizing Logger...");
29
         if Is Open (L.Logfile) then
30
            Close (L.Logfile);
31
         end if;
32
      end Finalize;
33
34
   end Loggers;
35
                                Listing 64: show_logger.adb
                          use Loggers;
   with Loggers;
1
   with Some_Processing;
2
3
   procedure Show_Logger is
4
      Name : aliased String := "report.log";
5
      Log : Logger (Name'Access);
6
   begin
      Put Line (Log, "Some info...");
8
      Some_Processing (Log);
```

9 end Show Logger; 10

1 2 3

4 5 6

7

8

9

10

11

12 13 14

15

16 17

18 19

21

22

23

24 25

26

27

28

7

Code block metadata

Project: Courses.Advanced_Ada.Resource_Management.Controlled_Types.Applications. →Logger MD5: d60ffbafd26d3d70a3d7807487dd95ab

Some info... Some processing...

Runtime output

```
Logger: Put_Line
Logger: Put_Line
Finalizing Logger...
```

This approach works, but requires us to declare an aliased string (Name), which we can give access to in the declaration of the Log object.

By encapsulating the file handling in controlled types, we ensure that files are properly opened when we want to use them, and that the files are closed when they're not going to be used anymore.