
Information technology — Object Management Group Object Constraint Language (OCL)

*Technologies de l'information — Langage de contraintes orienté-objet
(OCL) de l'OMG*



Reference number
ISO/IEC 19507:2012(E)

AVA Documents
Document Sharing Hub
© ISO/IEC 2012



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2012

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Table of Contents

Foreword	ix
Introduction	x
1 Scope	1
2 Conformance	1
3 References	2
3.1 Normative References	2
3.2 Informative References	2
4 Terms and Definitions	3
5 Notational Conventions	3
6 Additional Information	3
6.1 Changes to Adopted OMG Specifications	3
6.2 Structure of the Specification	3
6.3 Acknowledgements	4
7 OCL Language Description	5
7.1 General	5
7.2 Why OCL?	5
7.2.1 Where to Use OCL	5
7.3 Introduction	6
7.3.1 Legend	6
7.3.2 Example Class Diagram	6
7.3.3 Character Set	7
7.4 Relation to the UML Metamodel	7
7.4.1 Self	7
7.4.2 Specifying the UML Context	7
7.4.3 Invariants	8
7.4.4 Pre- and Postconditions	8
7.4.5 Package Context	9
7.4.6 Operation Body Expression	9
7.4.7 Initial and Derived Values	9
7.4.8 Other Types of Expressions	10
7.5 Basic Values and Types	10
7.5.1 Types from the UML Model	11
7.5.2 Enumeration Types	11
7.5.3 Let Expressions	11
7.5.4 Additional operations/attributes through «definition» expressions	12
7.5.5 Type Conformance	12

7.5.6 Re-typing or Casting	13
7.5.7 Precedence Rules	14
7.5.8 Use of Infix Operators	14
7.5.9 Keywords	15
7.5.10 Comment	16
7.5.11 Invalid Values	16
7.6 Objects and Properties	16
7.6.1 Properties: Attributes	17
7.6.2 Properties: Operations	17
7.6.3 Properties: AssociationEnds and Navigation	18
7.6.4 Navigation to Association Classes	20
7.6.5 Navigation from Association Classes	21
7.6.6 Navigation through Qualified Associations	21
7.6.7 Using Pathnames for Packages	21
7.6.8 Accessing overridden properties of supertypes	22
7.6.9 Predefined properties on All Objects	22
7.6.10 Features on Classes Themselves	23
7.6.11 Collections	24
7.6.12 Collections of Collections	25
7.6.13 Collection Type Hierarchy and Type Conformance Rules	25
7.6.14 Previous Values in Postconditions	25
7.6.15 Tuples	26
7.7 Collection Operations	27
7.7.1 Select and Reject Operations	27
7.7.2 Collect Operation	28
7.7.3 ForAll Operation	29
7.7.4 Exists Operation	30
7.7.5 Closure Operation	30
7.7.6 Iterate Operation	31
7.8 Messages in OCL	32
7.8.1 Calling operations and sending signals	32
7.8.2 Accessing result values	33
7.8.3 An example	33
7.9 Resolving Properties	34
8 Abstract Syntax	35
8.1 Introduction	35
8.2 The Types Package	35
8.2.1 Type Conformance	38
8.2.2 Operations and Well-formedness Rules for the Types Package	40
8.3 The Expressions Package	42
8.3.1 Expressions Core	43
8.3.2 FeatureCall Expressions	45
8.3.3 If Expressions	47
8.3.4 Message Expressions	48
8.4 Literal Expressions	49
8.4.1 Let Expressions	52
8.4.2 Well-formedness Rules of the Expressions package	53
8.4.3 Additional Operations on UML metaclasses	60

8.4.4 Additional Operations on OCL Metaclasses	62
8.4.5 Overview of class hierarchy of OCL Abstract Syntax metamodel	64
9 Concrete Syntax	65
9.1 General	65
9.2 Structure of the Concrete Syntax	65
9.3 A Note to Tool Builders	67
9.3.1 Parsing	67
9.3.2 Visibility	67
9.4 Concrete Syntax	67
9.4.1 ExpressionInOclCS	68
9.4.2 OclExpressionCS	68
9.4.3 VariableExpCS	69
9.4.4 simpleNameCS	69
9.4.5 restrictedKeywordCS	70
9.4.6 unreservedSimpleNameCS	71
9.4.7 pathNameCS	71
9.4.8 LiteralExpCS	72
9.4.9 EnumLiteralExpCS	72
9.4.10 CollectionLiteralExpCS	73
9.4.11 CollectionTypeIdentifierCS	73
9.4.12 CollectionLiteralPartsCS	74
9.4.13 CollectionLiteralPartCS	74
9.4.14 CollectionRangeCS	74
9.4.15 PrimitiveLiteralExpCS	75
9.4.16 TupleLiteralExpCS	76
9.4.17 UnlimitedNaturalLiteralExpCS	76
9.4.18 IntegerLiteralExpCS	76
9.4.19 RealLiteralExpCS	77
9.4.20 StringLiteralExpCS	77
9.4.21 BooleanLiteralExpCS	78
9.4.22 TypeLiteralExpCS	78
9.4.23 CallExpCS	79
9.4.24 LoopExpCS	79
9.4.25 IteratorExpCS	80
9.4.26 IterateExpCS	83
9.4.27 VariableDeclarationCS	84
9.4.28 TypeCS	85
9.4.29 primitiveTypeCS	85
9.4.30 oclTypeCS	86
9.4.31 collectionTypeCS	86
9.4.32 tupleTypeCS	87
9.4.33 variableDeclarationListCS	87
9.4.34 FeatureCallExpCS	88
9.4.35 OperationCallExpCS	88
9.4.36 PropertyCallExpCS	91
9.4.37 NavigationCallExpCS	93
9.4.38 AssociationClassCallExpCS	93
9.4.39 isMarkedPreCS	94
9.4.40 argumentsCS	94

9.4.41 LetExpCS	95
9.4.42 LetExpSubCS	95
9.4.43 OclMessageExpCS	96
9.4.44 OclMessageArgumentsCS	97
9.4.45 OclMessageArgCS	97
9.4.46 IfExpCS	98
9.4.47 NullLiteralExpCS	98
9.4.48 InvalidLiteralExpCS	99
9.4.49 Comments	99
9.5 Environment Definition	99
9.5.1 Environment	99
9.5.2 NamedElement	102
9.5.3 Namespace	102
9.6 Concrete to Abstract Syntax Mapping	102
9.7 Abstract Syntax to Concrete Syntax Mapping	103
10 Semantics Described Using UML	105
10.1 Introduction	105
10.2 The Values Package	106
10.2.1 Definitions of Concepts for the Values Package	107
10.2.2 Well-formedness Rules for the Values Package	111
10.2.3 Additional Operations for the Values Package	113
10.2.4 Overview of the Values Package	114
10.3 The Evaluations Package	115
10.3.1 Definitions of Concepts for the Evaluations Package	116
10.3.2 Well-formedness Rules of the Evaluations Package	125
10.3.3 Additional Operations of the Evaluations Package	132
10.3.4 Overview of the Values Package	133
10.4 The AS-Domain-Mapping Package	134
10.4.1 Well-formedness rules for the AS-Domain-Mapping.type-value Package	135
10.4.2 Additional Operations for the AS-Domain-Mapping.type-value Package	137
10.4.3 Well-formedness rules for the AS-Domain-Mapping.exp-eval Package	137
11 OCL Standard Library	145
11.1 Introduction	145
11.2 The OclAny, OclVoid, OclInvalid, and OclMessage Types	146
11.2.1 OclAny	146
11.2.2 OclMessage	146
11.2.3 OclVoid	146
11.2.4 OclInvalid	146
11.3 Operations and Well-formedness Rules	146
11.3.1 OclAny	146
11.3.2 OclVoid	148
11.3.3 OclMessage	148
11.4 Primitive Types	148
11.4.1 Real	148
11.4.2 Integer	148
11.4.3 String	149

11.4.4 Boolean	149
11.4.5 UnlimitedNatural	149
11.5 Operations and Well-formedness Rules	149
11.5.1 Real	149
11.5.2 Integer	150
11.5.3 String	151
11.5.4 Boolean	153
11.5.5 UnlimitedNatural	154
11.6 Collection-Related Types	155
11.6.1 Collection	156
11.6.2 Set	156
11.6.3 OrderedSet	156
11.6.4 Bag	156
11.6.5 Sequence	156
11.7 Operations and Well-formedness Rules	156
11.7.1 Collection	156
11.7.2 Set	159
11.7.3 OrderedSet	161
11.7.4 Bag	163
11.7.5 Sequence	165
11.8 Predefined Iterator Expressions	168
11.8.1 Extending the Standard Library with Iterator Expressions	168
11.9 Mapping Rules for Predefined Iterator Expressions	168
11.9.1 Collection	168
11.9.2 Set	170
11.9.3 Bag	171
11.9.4 Sequence	172
11.9.5 OrderedSet	173
12 The Use of OCL Expressions in UML Models	175
12.1 Introduction	175
12.2 The ExpressionInOcl Type	175
12.2.1 ExpressionInOcl	176
12.3 Well-formedness Rules	176
12.3.1 ExpressionInOcl	176
12.4 Standard Placements of OCL Expressions	177
12.4.1 How to Extend the Use of OCL at Other Places	177
12.5 Definition	177
12.5.1 Well-formedness Rules	177
12.6 Invariant	178
12.6.1 Well-formedness rules	178
12.7 Precondition	179
12.7.1 Well-formedness rules	179
12.7.2 Postcondition	180
12.7.3 Well-formedness rules	180
12.8 Initial Value Expression	180
12.8.1 Well-formedness rules	181
12.9 Derived Value Expression	182

12.10	Operation Body Expression	182
12.11	Guard	182
12.11.1	Well-formedness rules	183
12.12	Concrete Syntax of Context Declarations	183
12.12.1	packageDeclarationCS	184
12.12.2	contextDeclarationCS	184
12.12.3	propertyContextDeclCS	184
12.12.4	initOrDerValueCS	184
12.12.5	classifierContextDeclCS	184
12.12.6	invOrDefCS	184
12.12.7	defExpressionCS	185
12.12.8	operationContextDeclCS	185
12.12.9	prePostOrBodyDeclCS	185
12.12.10	operationCS	185
12.12.11	parametersCS	185
13	The Basic OCL and Essential OCL	187
13.1	Introduction	187
13.2	OCL Adaptation for Metamodeling	187
13.3	Diagrams	188
Annex A	Semantics	193
Annex B	Bibliography	229
Annex C	Legal Information	231

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 19507 was prepared by Technical Committee ISO/IEC JTC1, Information technology, in collaboration with the Object Management Group (OMG), following the submission and processing as a Publicly Available Specification (PAS) of the OMG Object Constraint Language specification Version 2.3.1.

ISO/IEC 19507, under the general title *Information technology - Open distributed processing - Object Constraint Language specification (OCL)*, apart from this introductory material is identical with that for the OMG specification for Object Constraint Language, v2.3.1.

Introduction

The rapid growth of distributed processing has led to a need for a coordinating framework for this standardization and ITU-T Recommendations X.901-904 | ISO/IEC 10746, the Reference Model of Open Distributed Processing (RM-ODP) provides such a framework. It defines an architecture within which support of distribution, interoperability, and portability can be integrated.

RM-ODP Part 2 (ISO/IEC 10746-2) defines the foundational concepts and modeling framework for describing distributed systems. The scopes and objectives of the RM-ODP Part 2 and the UML, while related, are not the same and, in a number of cases, the RM-ODP Part 2 and the UML specification use the same term for concepts that are related but not identical (e.g., interface). Nevertheless, a specification using the Part 2 modeling concepts can be expressed using UML with appropriate extensions (using stereotypes, tags, and constraints).

RM-ODP Part 3 (ISO/IEC 10746-3) specifies a generic architecture of open distributed systems, expressed using the foundational concepts and framework defined in Part 2. Given the relation between UML as a modeling language and Part 2 of the RM ODP standard, it is easy to show that UML is suitable as a notation for the individual viewpoint specifications defined by the RM-ODP.

OCL Language

OCL is a pure specification language; therefore, an OCL expression is guaranteed to be without side effects. When an OCL expression is evaluated, it simply returns a value. It cannot change anything in the model. This means that the state of the system will never change because of the evaluation of an OCL expression, even though an OCL expression can be used to *specify* a state change (e.g., in a post-condition).

OCL is not a programming language; therefore, it is not possible to write program logic or flow control in OCL. You cannot invoke processes or activate non-query operations within OCL. Because OCL is a modeling language in the first place, OCL expressions are not by definition directly executable.

OCL is a typed language so that each OCL expression has a type. To be well formed, an OCL expression must conform to the type conformance rules of the language. For example, you cannot compare an Integer with a String. Each Classifier defined within a UML model represents a distinct OCL type. In addition, OCL includes a set of supplementary predefined types. These are described in Clause 11 (“The OCL Standard Library”).

Information technology - Object Management Group Object Constraint Language (OCL)

1 Scope

This International Standard defines the Object Constraint Language (OCL), version 2.3.1. OCL version 2.3.1 is the version of OCL that is aligned with UML 2.3 and MOF 2.0.

2 Conformance

The UML 2.0 Infrastructure and the MOF 2.0 Core specifications that were developed in parallel with this OCL 2.3.1 specification share a common core. The OCL specification contains a well-defined and named subset of OCL that is defined purely based on the common core of UML and MOF. This allows this subset of OCL to be used with both the MOF and the UML, while the full specification can be used with the UML only.

The following compliance points are distinguished for both parts.

1. Syntax compliance: The tool can read and write OCL expressions in accordance with the grammar, including validating its type conformance and conformance of well-formedness rules against a model.
2. XMI compliance: The tool can exchange OCL expressions using XMI.
3. Evaluation compliance: The tool evaluates OCL expressions in accordance with the semantics clause. The following additional compliance points are optional for OCL evaluators, as they are dependent on the technical platform on which they are evaluated:
 - allInstances()
 - pre-values and oclIsNew() in postconditions
 - OclMessage
 - navigating across non-navigable associations
 - accessing private and protected features of an object

The following table shows the possible compliance points. Each tool is expected to fill in this table to specify which compliance points are supported.

Table 2.1 - Overview of OCL Compliance Points

	OCL-MOF subset	Full OCL
Syntax		
XMI		
Evaluation		
- allInstances		
- @pre in postconditions		
- OclMessage		
- navigating non-navigable associations		
- accessing private and protected features		

3 References

3.1 Normative References

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

- ISO 639 Codes for the representation of names of languages
- ISO 3166 Codes for the representation of names of countries and their subdivisions
- ISO/IEC 10646:2003 Information technology -- Universal Multiple-Octet Coded Character Set (UCS)
- UML 2.3 Superstructure Specification: <http://www.omg.org/spec/UML/2.3/Superstructure/PDF/>
- UML 2.3 Infrastructure Specification: <http://www.omg.org/spec/UML/2.3/Infrastructure/PDF/>
- MOF 2.0 Core Specification: <http://www.omg.org/spec/MOF/2.0/PDF/>
- UNICODE 5.1 Standard: <http://www.unicode.org/versions/Unicode5.1.0/>
- Unicode Technical Standard#10: <http://www.unicode.org/reports/tr10/>

3.2 Informative References

The following specifications are referenced in informative text:

- ISO/IEC 19501:2005 Information technology - Open Distributed Processing -- Unified Modeling Language (UML) Version 1.4.2

4 Terms and Definitions

There are no formal definitions that are taken from other documents.

5 Notational Conventions

There are no symbols defined.

6 Additional Information

6.1 Changes to Adopted OMG Specifications

This International Standard replaces the specification of OCL given in OCL 2.2.

The version of OCL specified in ISO/IEC 19501:2005 is intended for use in models based on UML 1.4.1 and UML 1.5. However, use of the OCL specified by ISO/IEC 19501:2005 is not prescribed by this specification.

The version of OCL specified in this International Standard is not directly applicable to models based on ISO/IEC 19501:2005.

6.2 Structure of the Specification

This International Standard is divided into several clauses.

- The OCL Language Description clause gives an informal description of OCL. This clause is not normative, but meant to be explanatory.
- Clause 8 (“Abstract Syntax”) describes the abstract syntax of OCL using a MOF 2.0 compliant metamodel. This is the same approach as used in the UML specifications. The metamodel is MOF compliant in the sense that it only uses constructs that are defined in the MOF.
- Clause 9 (“Concrete Syntax”) describes the canonical concrete syntax using an attributed EBNF grammar. This syntax is mapped onto the abstract syntax, achieving a complete separation between concrete and abstract syntax.
- Clause 10 (“Semantics Described using UML”) describes the semantics for OCL using UML.
- In Clause 11 (“The OCL Standard Library”) the OCL Standard Library is described. This defines type like Integer, Boolean, etc. and all the collection types. OCL is not a stand-alone language, but an integral part of the UML. An OCL expression needs to be placed within the context of a UML model.
- Clause 12 (“The Use of Ocl Expressions in UML Models”) describes a number of places within the UML where OCL expressions can be used.
- Clause 13 (“Basic OCL and Essential OCL”) defines the adaptation of the OCL metamodel when used in particular context of Core::Basic infrastructure library package and in the context of EMOF.
- Annex A (Semantics), Annex B (Bibliography), and Annex C (Legal Information)

6.3 Acknowledgements

The following companies submitted and/or supported parts of this specification:

- Adaptive Ltd.
- BoldSoft
- Borland Software Corporation
- Compuware Corporation
- Dresden University of Technology
- France Telecom
- International Business Machines
- IONA
- Kabira Technologies Inc.
- Kings College
- Classe Objecten
- Open Canarias, SL
- Oracle
- Project Technology Inc.
- Rational Software Corporation
- SAP AG
- Softeam
- Syntropy Ltd.
- Telelogic
- Thales
- University of Bremen
- University of Kent
- University of York
- Zeligsoft, Inc.

7 OCL Language Description

7.1 General

This clause introduces the Object Constraint Language (OCL), a formal language used to describe expressions on UML models. These expressions typically specify invariant conditions that must hold for the system being modeled or queries over objects described in a model. Note that when the OCL expressions are evaluated, they do not have side effects (i.e., their evaluation cannot alter the state of the corresponding executing system).

OCL expressions can be used to specify operations / actions that, when executed, do alter the state of the system. UML modelers can use OCL to specify application-specific constraints in their models. UML modelers can also use OCL to specify queries on the UML model, which are completely programming language independent.

Note - This clause is informative only and not normative.

7.2 Why OCL?

A UML diagram, such as a class diagram, is typically not refined enough to provide all the relevant aspects of a specification. There is, among other things, a need to describe additional constraints about the objects in the model. Such constraints are often described in natural language. Practice has shown that this will always result in ambiguities. In order to write unambiguous constraints, so-called formal languages have been developed. The disadvantage of traditional formal languages is that they are usable to persons with a strong mathematical background, but difficult for the average business or system modeler to use.

OCL has been developed to fill this gap. It is a formal language that remains easy to read and write. It has been developed as a business modeling language within the IBM Insurance division, and has its roots in the Syntropy method.

OCL is a pure specification language; therefore, an OCL expression is guaranteed to be without side effects. When an OCL expression is evaluated, it simply returns a value. It cannot change anything in the model. This means that the state of the system will never change because of the evaluation of an OCL expression, even though an OCL expression can be used to *specify* a state change (e.g., in a post-condition).

OCL is not a programming language; therefore, it is not possible to write program logic or flow control in OCL. You cannot invoke processes or activate non-query operations within OCL. Because OCL is a modeling language in the first place, OCL expressions are not by definition directly executable.

OCL is a typed language so that each OCL expression has a type. To be well formed, an OCL expression must conform to the type conformance rules of the language. For example, you cannot compare an Integer with a String. Each Classifier defined within a UML model represents a distinct OCL type. In addition, OCL includes a set of supplementary predefined types. These are described in Clause 11 ("The OCL Standard Library").

As a specification language, all implementation issues are out of scope and cannot be expressed in OCL.

The evaluation of an OCL expression is instantaneous. This means that the states of objects in a model cannot change during evaluation.

7.2.1 Where to Use OCL

OCL can be used for a number of different purposes:

- as a query language,

- to specify invariants on classes and types in the class model,
- to specify type invariant for Stereotypes,
- to describe pre- and post conditions on Operations and Methods,
- to describe Guards,
- to specify target (sets) for messages and actions,
- to specify constraints on operations, and
- to specify derivation rules for attributes for any expression over a UML model.

7.3 Introduction

7.3.1 Legend

Text written in the typeface as shown below is an OCL expression.

'This is an OCL expression'

The *context* keyword introduces the context for the expression. The keyword *inv*, *pre*, and *post* denote the stereotypes, respectively «invariant», «precondition», and «postcondition» of the constraint. The actual OCL expression comes after the colon.

context TypeName **inv**:

'this is an OCL expression with stereotype <<invariant>> in the
context of TypeName' = 'another string'

In the examples the keywords of OCL are written in boldface. The boldface has no formal meaning, but is used to make the expressions more readable. OCL expressions are written using ASCII characters only.

Words in *Italics* within the main text of the paragraphs refer to parts of OCL expressions.

7.3.2 Example Class Diagram

The diagram below is used in the examples in this clause.




Figure 7.1 - Class Diagram Example

7.3.3 Character Set

OCL text comprises characters in the UNICODE character set. In particular, string literals, comments, and the names of types, features, and other elements in the UML model may contain any valid UNICODE character.

7.4 Relation to the UML Metamodel

7.4.1 Self

Each OCL expression is written in the context of an instance of a specific type. In an OCL expression, the reserved word *self* is used to refer to the contextual instance. For example, if the context is *Company*, then *self* refers to an instance of *Company*.

7.4.2 Specifying the UML Context

The context of an OCL expression within a UML model can be specified through a so-called context declaration at the beginning of an OCL expression. The context declaration of the constraints in the following sub clauses is shown.

If the constraint is shown in a diagram, with the proper stereotype and the dashed lines to connect it to its contextual element, there is no need for an explicit context declaration in the test of the constraint. The context declaration is optional.

7.4.3 Invariants

The OCL expression can be part of an Invariant, which is a Constraint stereotyped as an «invariant». When the invariant is associated with a Classifier, the latter is referred to as a “type” in this clause. An OCL expression is an invariant of the type and must be true for all instances of that type at any time. (Note that all OCL expressions that express invariants are of the type Boolean.)

For example, if in the context of the Company type in Figure 7.1, the following expression would specify an invariant that the number of employees must always exceed 50:

```
self.numberOfEmployees > 50
```

where *self* is an instance of type Company. (We can view *self* as the object from where we start evaluating the expression.) This invariant holds for every instance of the Company type.

The type of the contextual instance of an OCL expression, which is part of an invariant, is written with the *context* keyword, followed by the name of the type as follows. The label *inv*: declares the constraint to be an «invariant» constraint.

```
context Company inv:  
  self.numberOfEmployees > 50
```

In most cases, the keyword *self* can be dropped because the context is clear, as in the above examples. As an alternative for *self*, a different name can be defined playing the part of *self*. For example:

```
context c : Company inv:  
  c.numberOfEmployees > 50
```

This invariant is equivalent to the previous one.

Optionally, the name of the constraint may be written after the *inv* keyword, allowing the constraint to be referenced by name. In the following example the name of the constraint is *enoughEmployees*.

```
context c : Company inv enoughEmployees:  
  c.numberOfEmployees > 50
```

7.4.4 Pre- and Postconditions

The OCL expression can be part of a Precondition or Postcondition, corresponding to «precondition» and «postcondition» stereotypes of Constraint associated with an Operation or other behavioral feature. The contextual instance *self* then is an instance of the type that owns the operation or method as a feature. The context declaration in OCL uses the *context* keyword, followed by the type and operation declaration. The stereotype of constraint is shown by putting the labels ‘pre:’ and ‘post:’ before the actual Preconditions and Postconditions. For example:

```
context TypeName::operationName(param1 : Type1, ... ) : ReturnType  
  pre : param1 > ...  
  post: result = ...
```

The name *self* can be used in the expression referring to the object on which the operation was called. The reserved word *result* denotes the result of the operation, if there is one. The names of the parameters (*param1*) can also be used in the OCL expression. In the example diagram, we can write:

```
context Person::income(d : Date) : Integer  
  post: result = 5000
```

Optionally, the name of the precondition or postcondition may be written after the *pre* or *post* keyword, allowing the constraint to be referenced by name. In the following example the name of the precondition is *parameterOk* and the name of the postcondition is *resultOk*. In the UML metamodel, these names are the values of the attribute *name* of the metaclass *Constraint* that is inherited from *ModelElement*.

```
context Typename::operationName(param1 : Type1, ... ): ReturnType
  pre parameterOk: param1 > ...
  post resultOk : result = ...
```

7.4.5 Package Context

The above context declaration is precise enough when the package in which the Classifier belongs is clear from the environment. To specify explicitly in which package invariant, pre or postcondition Constraints belong, these constraints can be enclosed between ‘package’ and ‘endpackage’ statements. The package statements have the syntax:

```
package Package::SubPackage

context X inv:
  ... some invariant ...
context X::operationName(...)
  pre: ... some precondition ...

endpackage
```

An OCL file (or stream) may contain any number package statements, thus allowing all invariant, preconditions, and postconditions to be written and stored in one file. This file may co-exist with a UML model as a separate entity.

7.4.6 Operation Body Expression

An OCL expression may be used to indicate the result of a query operation. This can be done using the following syntax:

```
context Typename::operationName(param1 : Type1, ... ): ReturnType
body: -- some expression
```

The expression must conform to the result type of the operation. Like in the pre- and postconditions, the parameters may be used in the expression. Pre-, and postconditions, and body expressions may be mixed together after one operation context. For example:

```
context Person::getCurrentSpouse() : Person
pre: self.isMarried = true
body: self.mariages->select( m | m.ended = false ).spouse
```

7.4.7 Initial and Derived Values

An OCL expression may be used to indicate the initial or derived value of an attribute or association end. This can be done using the following syntax:

```
context Typename::attributeName: Type
init: -- some expression representing the initial value

context Typename::assocRoleName: Type
derive: -- some expression representing the derivation rule
```

The expression must conform to the result type of the attribute. In the case the context is an association end the expression must conform to the classifier at that end when the multiplicity is at most one, or Set, or OrderedSet when the multiplicity may be more than one. Initial and derivation expressions may be mixed together after one context. For example:

```
context Person::income : Integer
init: parents.income->sum() * 1% -- pocket allowance
derive: if underAge
        then parents.income->sum() * 1% -- pocket allowance
        else job.salary          -- income from regular job
endif
```

The derivation constraint must be satisfied at any time, hence the derivation includes the initialization. Both are allowed on the same property but they must not be contradictory. For each property there should be at most one initialization constraint and at most one derivation constraint.

7.4.8 Other Types of Expressions

Any OCL expression can be used as the value for an attribute of the UML metaclass Expression or one of its subtypes. In that case, the semantics sub clause describes the meaning of the expression. A special subclass of Expression, called ExpressionInOcl is used for this purpose. See Clause 12 “The Use of OCL Expressions in UML Models” for a definition.

7.5 Basic Values and Types

In OCL, a number of basic types are predefined and available to the modeler at all times. These predefined value types are independent of any object model and are part of the definition of OCL.

The most basic value in OCL is a value of one of the basic types. The basic types of OCL, with corresponding examples of their values, are shown in the following table.

Table 7.1- Basic OCL types and their values

type	values	consistent with implementation definitions
OclInvalid	invalid	
OclVoid	null, invalid	
Boolean	true, false	(MOF) http://www.w3.org/TR/xmlschema-2/#boolean
Integer	1, -5, 2, 34, 26524, ...	(MOF) http://www.w3.org/TR/xmlschema-2/#integer
Real	1.5, 3.14, ...	http://www.w3.org/TR/xmlschema-2/#double
String	'To be or not to be...'	(MOF) http://www.w3.org/TR/xmlschema-2/#string
UnlimitedNatural	0, 1, 2, 42, ..., *	http://www.w3.org/TR/xmlschema-2/#nonNegativeInteger

OCL defines a number of operations on the predefined types. Table 7.2 gives some examples of the operations on the predefined types. See 11.4, 'Primitive Types' for a complete list of all operations.

Table 7.2- Examples of operations on the predefined types

type	operations
Integer	*, +, -, /, abs()
Real	*, +, -, /, floor()

Table 7.2- Examples of operations on the predefined types

Boolean	and, or, xor, not, implies, if-then-else
String	concat(), size(), substring()
UnlimitedNatural	*, +, /

Collection, Set, Bag, Sequence, and Tuple are basic types as well. Their specifics will be described in the upcoming sub clauses.

Multiple adjacent strings are concatenated allowing a long string to be specified on multiple lines.

```
'This is a '
'concatenated string'      -- 'This is a concatenated string'
```

Unicode characters are used within single quoted sequences, with the following backslash based escape sequences used to define backslash and other characters.

```
\b      -- backspace
\t      -- horizontal tab
\n      -- linefeed
\f      -- form feed
\r      -- carriage return
\"      -- double quote
\'      -- single quote
\\      -- backslash
\xhh    -- #x00 to #xFF
\uhhhh  -- #x0000 to #xFFFF
```

where *h* is a hex digit: 0 to 9, A to F or a to f.

Reserved words such as true and arbitrary awkward spellings may be used as names by enclosing the name in underscore-prefixed single quotes.

```
self.'if' = 'tabbed\tvariable'. 'spaced operation'()
```

7.5.1 Types from the UML Model

Each OCL expression is written in the context of a UML model, a number of classifiers (types/classes, ...), their features and associations, and their generalizations. All classifiers from the UML model are types in the OCL expressions that are attached to the model.

7.5.2 Enumeration Types

Enumerations are Datatypes in UML and have a name, just like any other Classifier. An enumeration defines a number of enumeration literals that are the possible values of the enumeration. Within OCL one can refer to the value of an enumeration. When we have Datatype named Gender in the example model with values 'female' or 'male' they can be used as follows:

```
context Person inv: gender = Gender::male
```

7.5.3 Let Expressions

Sometimes a sub-expression is used more than once in a constraint. The *let* expression allows one to define a variable that can be used in the constraint.

```
context Person inv:
```

```

let income : Integer = self.job.salary->sum() in
if isUnemployed then
    income < 100
else
    income >= 100
endif

```

A let expression may be included in any kind of OCL expression. It is only known within this specific expression. A variable declaration inside a let must have a declared type and an initial value.

7.5.4 Additional operations/attributes through «definition» expressions

The Let expression allows a variable to be used in one Ocl expression. To enable reuse of variables/operations over multiple OCL expressions one can use a Constraint with the stereotype «definition», in which helper variables/operations are defined. This «definition» Constraint must be attached to a Classifier and may only contain variable and/or operation definitions, nothing else. All variables and operations defined in the «definition» constraint are known in the same context as where any property of the Classifier can be used. Such variables and operations are attributes and operations with stereotype «OclHelper» of the classifier. They are used in an OCL expression in exactly the same way as normal attributes or operations are used. The syntax of the attribute or operation definitions is similar to the Let expression, but each attribute and operation definition is prefixed with the keyword 'def' as shown below.

```

context Person
def: income : Integer = self.job.salary->sum()
def: nickname : String = 'Little Red Rooster'
def: hasTitle(t : String) : Boolean = self.job->exists(title = t)

```

Operations or attributes defined by "definitions expressions" may be static (classifier scoped). In that case the static keyword should be used before "def."

```

context MyClass
static def: globalId() : Integer = ...

```

The names of the attributes / operations in a let expression may not conflict with the names of respective attributes/ associationEnds and operations of the Classifier.

Using this definition syntax is identical to defining an attribute/operation in the UML with stereotype «OclHelper» with an attached OCL constraint for its derivation.

7.5.5 Type Conformance

OCL is a typed language and the basic value types are organized in a type hierarchy. This hierarchy determines conformance of the different types to each other. You cannot, for example, compare an Integer with a Boolean or a String.

An OCL expression in which all the types conform is a valid expression. An OCL expression in which the types don't conform is an invalid expression. It contains a *type conformance error*. A type *type1* conforms to a type *type2* when an instance of *type1* can be substituted at each place where an instance of *type2* is expected. The type conformance rules for types in the class diagrams are simple.

- Each type conforms to each of its supertypes.
- Type conformance is transitive: if *type1* conforms to *type2*, and *type2* conforms to *type3*, then *type1* conforms to *type3*.

The effect of this is that a type conforms to its supertype, and all the supertypes above. The type conformance rules for the types from the OCL Standard Library are listed in Table 7.3., where the third column specifies an additional condition

Table 7.3- Type conformance rules

Type	Conforms to/Is a subtype of	Condition
Set(T1)	Collection(T2)	if T1 conforms to T2
Sequence(T1)	Collection(T2)	if T1 conforms to T2
Bag(T1)	Collection(T2)	if T1 conforms to T2
OrderedSet(T1)	Collection(T2)	if T1 conforms to T2
Integer	Real	
UnlimitedNatural	Integer	* is an <i>invalid</i> Integer

which must be satisfied by the involved types to verify the type conformance rule

Although UnlimitedNatural conforms to Integer, '*' is an invalid Integer, so that the evaluation of the expression '1 + *' results in invalid.

The conformance relation between the collection types only holds if they are collections of element types that conform to each other. See 7.6.13, 'Collection Type Hierarchy and Type Conformance Rules' for the complete conformance rules for collections.

Table 7.4 provides examples of valid and invalid expressions.

Table 7.4- Valid and Invalid Expressions

OCL expression	valid	explanation
1 + 2 * 34	yes	
1 + 'motorcycle'	no	type String does not conform to type Integer
23 * false	no	type Boolean does not conform to Integer
12 + 13.5	yes	

7.5.6 Re-typing or Casting

In some circumstances, it is desirable to use a property of an object that is defined on a subtype of the current known type of the object. Because the property is not defined on the current known type, this results in a type conformance error.

When it is certain that the actual type of the object is the subtype, the object can be re-typed using the operation *oclAsType(Classifier)*. This operation results in the same object, but the known type is the argument *Classifier*. When there is an object *object* of type *Type1* and *Type2* is another type, it is allowed to write:

object.oclAsType(Type2) --- changes the static type of the expression to Type2

An object can only be re-typed to a type to which it conforms. If the actual type of the object, at evaluation time, is not a subtype of the type to which it is re-typed, then the result of *oclAsType* is *invalid*.

Casting provides visibility, at parse time, of features not defined in the context of an expression's static type. It does not coerce objects to instances of another type, nor can it provide access to hidden or overridden features of a type. For this, the feature call is qualified by the name of the type (a path name, if necessary) whose definition of the feature is to be accessed.

For example, if class *Employee* redefines the *age()* : *Integer* operation of the *Person* class, a constraint may access the *Person* definition as in

context Employee

```
inv: self.age() <= self.Person::age()
```

For clarity, the qualified form may only be used with an explicit source expression.

7.5.7 Precedence Rules

The precedence order for the operations, starting with highest precedence, in OCL is:

- literal and variable expressions, “(“ and “)”, “if-then-else-endif”
- “let-in”
- @pre
- call expressions: “^”, “^^”, “.” and “->”
- unary “not” and unary “-”
- “*” and “/”
- “+” and binary “-”
- “<”, “>”, “<=”, “>=”
- “=”, “<>”
- “and”
- “or”
- “xor”
- “implies”
- “in”

All infix operators are left associative, equal precedence operators are evaluated left to right.

A let expression is both high precedence and low precedence; high on the left so that a let expression behaves as an atomic value in operations, low on the right so that the in-expression can be an arbitrary expression. “a + let ... in a + let ... in a + a” is “a + (let ... in (a + (let ... in (a + a)))).”

Parentheses “(“ and “)” can be used to change precedence and associativity.

7.5.8 Use of Infix Operators

The use of infix operators is allowed in OCL. The operators ‘+,’ ‘-,’ ‘*,’ ‘/,’ ‘<,’ ‘>,’ ‘<>,’ ‘<=,’ ‘>=,’ are used as infix operators. If a type defines one of those operators with the correct signature, they will be used as infix operators. The expression:

$$a + b$$

is equal to the expression:

a. '+'(b)

that is, invoking the “+” operation on a with b as the parameter to the operation.

The infix operators defined for a type must have exactly one parameter. For the infix operators ‘<,’ ‘>,’ ‘<=,’ ‘>=,’ ‘<>,’ ‘and,’ ‘or,’ and ‘xor’ the return type must be Boolean.

7.5.9 Keywords

Keywords in OCL are reserved words. That means that the keywords cannot occur as a name. A reserved word may be used as the name of a package, a type, a feature, a variable or a constraint by enclosing the word in underscore-prefixed single quotes. The list of keywords is shown below:

- and
- body
- context
- def
- derive
- else
- endif
- endpackage
- false
- if
- implies
- in
- init
- inv
- invalid
- let
- not
- null
- or
- package
- post
- pre
- self
- static
- then
- true
- xor

The following words are restricted. A restricted word can only be used as a name when preceded by a “::”. A restricted word may also be used by enclosing the word in underscore-prefixed single quotes.

- Bag
- Boolean
- Collection
- Integer
- OclAny
- OclInvalid
- OclMessage
- OclVoid
- OrderedSet
- Real
- Sequence
- Set

String
Tuple
UnlimitedNatural

Note that operation names such as `iterate`, `forAll`, and `oclType` are not reserved or restricted.

7.5.10 Comment

Comments in OCL are written following two successive dashes (minus signs). Everything immediately following the two dashes up to and including the end of line is part of the comment.

For example:

```
-- this is a comment
```

7.5.11 Invalid Values

Some expressions will, when evaluated, have an invalid value. For instance, typecasting with `oclAsType()` to a type that the object does not support or getting the `->first()` element of an empty collection will result in invalid. In general, an expression where one of the parts is null or invalid will itself be invalid. There are some important exceptions to this rule, however. First, there are the logical operators:

- True OR-ed with anything is True
- False AND-ed with anything is False
- False IMPLIES anything is True
- anything IMPLIES True is True

The rules for OR and AND are valid irrespective of the order of the arguments and they are valid whether the value of the other sub-expression is known or not.

The IF-expression is another exception. It will be valid as long as the chosen branch is valid, irrespective of the value of the other branch.

Finally, there is an explicit operation for testing if the value of an expression is undefined. `oclIsUndefined()` is an operation on `OclAny` that results in True if its argument is *null* or *invalid* and False otherwise.

7.6 Objects and Properties

OCL expressions can refer to Classifiers, e.g., types, classes, interfaces, associations (acting as types), and datatypes. Also all attributes, association-ends, methods, and operations without side effects that are defined on these types, etc. can be used. In a class model, an operation or method is defined to be side effect free if the `isQuery` attribute of the operations is true. For the purpose of this document, we will refer to attributes, association-ends, and side effect free methods and operations as being *properties*. A property is one of:

- an Attribute
- an AssociationEnd
- an Operation with *isQuery* being true
- a Method with *isQuery* being true

The value of a property on an object that is defined in a class diagram is specified in an OCL expression by a dot followed by the name of the property. For example:

```
context Person inv:
    self.isMarried
```

If *self* is a reference to an object, then *self.property* is the value of the *property* property on *self*.

7.6.1 Properties: Attributes

For example, the age of a Person is written as *self.age*:

```
context Person inv:
    self.age > 0
```

The value of the subexpression *self.age* is the value of the *age* attribute on the particular instance of Person identified by *self*. The type of this subexpression is the type of the attribute *age*, which is the standard type Integer.

Using attributes and operations defined on the basic value types, we can express calculations etc. over the class model. For example, a business rule might be “the age of a Person is always greater than zero.” This can be stated by the invariant above.

Attributes may have multiplicities in a UML model. Whenever the multiplicity of an attribute is greater than 1, the result type is collection of values. Collections in OCL are described later in this clause.

7.6.2 Properties: Operations

Operations may have parameters. For example, as shown earlier, a Person object has an income expressed as a function of the date. This operation would be accessed as follows, for a Person *aPerson* and a date *aDate*:

```
aPerson.income(aDate)
```

The result of this operation call is a value of the return type of the operation, which is Integer in this example. If the operation has out or in/out parameters, the result of this operation is a tuple containing all out, in/out parameters and the return value. For example, if the income operation would have an out parameter *bonus*, the result of the above operation call is of type *Tuple(bonus: Integer, result: Integer)*. You can access these values using the names of the out parameters, and the keyword *result*. For example:

```
aPerson.income(aDate).bonus = 300 and
aPerson.income(aDate).result = 5000
```

Note that the out parameters need not be included in the operation call. Values for all in or in/out parameters are necessary.

Defining operations

The operation itself could be defined by a postcondition constraint. This is a constraint that is stereotyped as «postcondition». The object that is returned by the operation can be referred to by *result*. It takes the following form:

```
context Person::income (d: Date) : Integer
    post: result = age * 1000
```

The right-hand-side of this definition may refer to the operation being defined (i.e., the definition may be recursive) as long as the recursion is not infinite. Inside a pre- or postcondition one can also use the parameters of the operation. The type of *result*, when the operation has no out or in/out parameters, is the return type of the operation, which is Integer in the above example. When the operation does have out or in/out parameters, the return type is a Tuple as explained above. The postcondition for the income operation with out parameter *bonus* may take the following form:

```

context Person::income (d: Date, bonus: Integer) : Integer
  post: result = Tuple { bonus = ...,
                      result = .... }

```

To refer to an operation or a method that doesn't take a parameter, parentheses with an empty argument list are mandatory:

```

context Company inv:
  self.stockPrice() > 0

```

7.6.3 Properties: AssociationEnds and Navigation

Starting from a specific object, we can navigate an association on the class diagram to refer to other objects and their properties. To do so, we navigate the association by using the opposite association-end:

```
object.associationEndName
```

The value of this expression is the set of objects on the other side of the *associationEndName* association. If the multiplicity of the association-end has a maximum of one ("0..1" or "1"), then the value of this expression is an object. In the example class diagram, when we start in the context of a Company (i.e., *self* is an instance of Company), we can write:

```

context Company
  inv: self.manager.isUnemployed = false
  inv: self.employee->notEmpty()

```

In the first invariant *self.manager* is a Person, because the multiplicity of the association is one. In the second invariant *self.employee* will evaluate in a Set of Persons. By default, navigation will result in a Set. When the association on the Class Diagram is adorned with {ordered}, the navigation results in an OrderedSet.

Collections, like Sets, OrderedSets, Bags, and Sequences are predefined types in OCL. They have a large number of predefined operations on them. A property of the collection itself is accessed by using an arrow '->' followed by the name of the property. The following example is in the context of a person:

```

context Person inv:
  self.employer->size() < 3

```

This applies the *size* property on the Set *self.employer*, which results in the number of employers of the Person *self*.

```

context Person inv:
  self.employer->isEmpty()

```

This applies the *isEmpty* property on the Set *self.employer*. This evaluates to true if the set of employers is empty and false otherwise.

Missing Association names

The association name is never missing. If no explicit name is available, an implicit name is constructed in accordance with the UML style guide. Associations that are not explicitly named, are given names that are constructed according to the following production rule:

```
"A_" <association-end-name1> "_" <association-end-name2>
```

where <association-end-name1> is the name of one association end and lexically precedes <association-end-name2> which is the name of the other association end.

Missing Association End names

The name of an association-end is never missing. If no explicit name is available an implicit name is taken from the name of the class to which the end is attached.

NOTE: To tool vendors: this is a non-normative change from OCL 2.2, where the UML style guidance of converting the first letter of the implicit name to lowercase was endorsed. The normative text has never defined how implicit names are obtained. Tool vendors may wish to provide backward/forward compatibility warnings for this change.




Figure 7.2 - Ambiguous name example

This may result in an ambiguity between an implicit association end name and another explicit name, unless only one of the association ends is navigable. The ambiguous name cannot be used in OCL.

`aPerson.role` -- ambiguous

Qualifying association ends with association names

An association end name may be qualified with its association name or its source classifier name to resolve an ambiguity.

<code>aPerson.Person::role</code>	-- still ambiguous
<code>aPerson.A_person_role::role</code>	-- some Parts, using implicit Person to Part association name
<code>aPerson.A_owner_role::role</code>	-- a Role, using implicit Person to Role association name

Ends owned by associations

In a UML association, an end may be owned by the Classifier at that end, or by the association, itself. The ownership of the end is not significant to OCL. In either case, the association end is considered as a property of the Classifier and can be navigated from that end to the other.

Navigation over Associations with Multiplicity Zero or One

Because the multiplicity of the role manager is one, *self.manager* is an object of type Person. Such a single object can be used as a Set as well. It then behaves as if it is a Set containing the single object. The usage as a set is done through the arrow followed by a property of Set. This is shown in the following example:

```

context Company inv:
  self.manager->size() = 1
  
```

The sub-expression *self.manager* is used as a Set, because the arrow is used to access the *size* property on Set. This expression evaluates to true.

```

context Company inv:
  self.manager->foo
  
```

The sub-expression *self.manager* is used as Set, because the arrow is used to access the *foo* property on the Set. This expression is incorrect, because *foo* is not a defined property of Set.

```

context Company inv:
  self.manager.age > 40
  
```

The sub-expression *self.manager* is used as a Person, because the dot is used to access the *age* property of Person.

In the case of an optional (0..1 multiplicity) association, this is especially useful to check whether there is an object or not when navigating the association. In the example we can write:

```
context Person inv:
    self.wife->notEmpty() implies self.wife.gender = Gender::female
```

Combining Properties

Properties can be combined to make more complicated expressions. An important rule is that an OCL expression always evaluates to a specific object of a specific type. After obtaining a result, one can always apply another property to the result to get a new result value. Therefore, each OCL expression can be read and evaluated left-to-right.

Following are some invariants that use combined properties on the example class diagram:

[1] [1] Married people are of age ≥ 18

```
context Person inv:
    self.wife->notEmpty() implies self.wife.age  $\geq 18$  and
    self.husband->notEmpty() implies self.husband.age  $\geq 18$ 
```

[2] [2] a company has at most 50 employees

```
context Company inv:
    self.employee->size()  $\leq 50$ 
```

7.6.4 Navigation to Association Classes

To specify navigation to association classes (Job and Marriage in the example), OCL uses a dot and the name of the association class:

```
context Person inv:
    self.Job
```

The sub-expression *self.Job* evaluates to a Set of all the jobs a person has with the companies that are his/her employer. In the case of an association class, there is no explicit rolename in the class diagram. The name *Job* used in this navigation is the name of the association class.

In case of a recursive association, that is an association of a class with itself, the name of the association class alone is not enough. We need to distinguish the direction in which the association is navigated as well as the name of the association class. Take the following model as an example.




Figure 7.3 - Navigating recursive association classes

When navigating to an association class such as *EmployeeRanking* there are two possibilities depending on the direction. For instance, in the above example, we may navigate towards the *employees* end, or the *bosses* end. By using the name of the association class alone, these two options cannot be distinguished. To make the distinction, the rolename of the direction in which we want to navigate is added to the association class name, enclosed in square brackets. In the expression

```

context Person inv:
    self.EmployeeRanking[bosses]->sum() > 0

```

the *self.EmployeeRanking[bosses]* evaluates to the set of *EmployeeRankings* belonging to the collection of *bosses*. And in the expression

```

context Person inv:
    self.EmployeeRanking[employees]->sum() > 0

```

the *self.EmployeeRanking[employees]* evaluates to the set of *EmployeeRankings* belonging to the collection of *employees*. The unqualified use of the association class name is not allowed in such a recursive situation. Thus, the following example is invalid:

```

context Person inv:
    self.EmployeeRanking->sum() > 0 -- INVALID!

```

In a non-recursive situation, the association class name alone is enough, although the qualified version is allowed as well. Therefore, the examples at the start of this sub clause could also be written as:

```

context Person inv:
    self.Job[employer]

```

7.6.5 Navigation from Association Classes

We can navigate from the association class itself to the objects that participate in the association. This is done using the dot-notation and the role-names at the association-ends.

```

context Job
    inv: self.employer.numberOfEmployees >= 1
    inv: self.employee.age > 21

```

Navigation from an association class to one of the objects on the association will always deliver exactly one object. This is a result of the definition of *AssociationClass*. Therefore, the result of this navigation is exactly one object, although it can be used as a *Set* using the arrow (*->*).

7.6.6 Navigation through Qualified Associations

Qualified associations use one or more qualifier attributes to select the objects at the other end of the association. To navigate them, we can add the values for the qualifiers to the navigation. This is done using square brackets, following the role-name. It is permissible to leave out the qualifier values, in which case the result will be all objects at the other end of the association. The following example results in a *Set(Person)* containing all customers of the Bank.

```

context Bank inv:
    self.customer

```

The next example results in one *Person*, having account number 8764423.

```

context Bank inv:
    self.customer[8764423]

```

If there is more than one qualifier attribute, the values are separated by commas, in the order which is specified in the UML class model. It is not permissible to partially specify the qualifier attribute values.

7.6.7 Using Pathnames for Packages

Within UML, types are organized in packages. OCL provides a way of explicitly referring to types in other packages by using a package-pathname prefix. The syntax is a package name, followed by a double colon:

PackageName::Typename

This usage of pathnames is transitive and can also be used for packages within packages:

PackageName1::PackageName2::Typename

7.6.8 Accessing overridden properties of supertypes

Whenever properties are redefined within a type, the property of the supertypes can be accessed using the *oclAsType()* operation. Whenever we have a class B as a subtype of class A, and a property p1 of both A and B, we can write:

```
context B inv:
  self.oclAsType(A).p1 -- accesses the p1 property defined in A
  self.p1 -- accesses the p1 property defined in B
```

Figure 7.4 shows an example where such a construct is needed. In this model fragment there is an ambiguity with the OCL expression on Dependency:

```
context Dependency inv:
  self.source <> self
```

This can either mean normal association navigation, which is inherited from ModelElement, or it might also mean navigation through the dotted line as an association class. Both possible navigations use the same role-name, so this is always ambiguous. Using *oclAsType()* we can distinguish between them with:

```
context Dependency
  inv: self.oclAsType(Dependency).source->isEmpty()
  inv: self.oclAsType(ModelElement).source->isEmpty()
```




Figure 7.4 - Accessing Overridden Properties Example

7.6.9 Predefined properties on All Objects

There are several properties that apply to all objects, and are predefined in OCL. These are:

```
oclIsTypeOf (t : Classifier) : Boolean
oclIsKindOf (t : Classifier) : Boolean
oclInState (s : OclState) : Boolean
oclIsNew () : Boolean
oclAsType (t : Classifier) : instance of Classifier
```

The operation *oclIsTypeOf* results in true if the *type* of self and *t* are the same. For example:

```
context Person
  inv: self.oclIsTypeOf( Person )    -- is true
  inv: self.oclIsTypeOf( Company )   -- is false
```

The above property deals with the direct type of an object. The *oclIsKindOf* property determines whether *t* is either the direct type or one of the supertypes of an object.

The operation *oclInState(s)* results in true if the object is in the state *s*. Possible states for the operation *oclInState(s)* are all states of the statemachine that defines the classifier's behavior. For nested states the statenames can be combined using the double colon "::".




Figure 7.5 - Statemachine Example

In the example statemachine above, values for *s* can be *On*, *Off*, *Off::Standby*, *Off::NoPower*. If the classifier of *object* has the above associated statemachine, valid OCL expressions are:

```
object.oclInState(On)
object.oclInState(Off)
object.oclInState(Off::Standby)
object.oclInState(Off::NoPower)
```

If there are multiple statemachines attached to the object's classifier, then the statename can be prefixed with the name of the statemachine containing the state and the double colon '::,' as with nested states.

The operation *oclIsNew* evaluates to true if, used in a postcondition, the object is created during performing the operation (i.e., it didn't exist at precondition time).

7.6.10 Features on Classes Themselves

All properties discussed until now in OCL are properties on instances of classes. The types are either predefined in OCL or defined in the class model. In OCL, it is also possible to use static features, applicable to the types/classes themselves rather than to their instances. For example, the *Employee* class may define a static operation "uniqueID" that computes a unique ID to use in the initialization of the employee ID attribute:

```
context Employee::id : String init:
  Employee::uniqueID()
```

Static features are invoked using the '::' operator and are distinct from the features of the Classifier metaclass, which include the *allInstances* operation pre-defined by OCL. If we want to make sure that all instances of *Person* have unique names, we can write:

```
context Person inv:
  Person.allInstances()->forAll(p1, p2 |
    p1 <> p2 implies p1.name <> p2.name)
```

Invocation of *allInstances* uses the '.' operator rather than '::' because it is not a static operation. It is an operation applicable to instances of the Classifier metaclass, of which Person is an example.

7.6.11 Collections

Single navigation of an association results in a Set, combined navigations in a Bag, and navigation over associations adorned with {ordered} results in an OrderedSet. Therefore, the collection types defined in the OCL Standard Library play an important role in OCL expressions.

The type Collection is predefined in OCL. The Collection type defines a large number of predefined operations to enable the OCL expression author (the modeler) to manipulate collections. Consistent with the definition of OCL as an expression language, collection operations never change collections; *isQuery* is always true. They may result in a collection, but rather than changing the original collection they project the result into a new one.

Collection is an abstract type, with the concrete collection types as its subtypes. OCL distinguishes three different collection types: Set, Sequence, and Bag. A Set is the mathematical set. It does not contain duplicate elements. A Bag is like a set, which may contain duplicates (i.e., the same element may be in a bag twice or more). A Sequence is like a Bag in which the elements are ordered. Both Bags and Sets have no order defined on them.

Collection Literals

Sets, Sequences, and Bags can be specified by a literal in OCL. Curly brackets surround the elements of the collection, elements in the collection are written within, separated by commas. The type of the collection is written before the curly brackets:

```
Set { 1 , 2 , 5 , 88 }
Set { 'apple,' 'orange,' 'strawberry' }
```

A Sequence:

```
Sequence { 1, 3, 45, 2, 3 }
Sequence { 'ape,' 'nut' }
```

A bag:

```
Bag {1, 3, 4, 3, 5 }
```

Because of the usefulness of a Sequence of consecutive Integers, there is a separate literal to create them. The elements inside the curly brackets can be replaced by an interval specification, which consists of two expressions of type Integer, *Int-expr1* and *Int-expr2*, separated by '..'. This denotes all the Integers between the values of *Int-expr1* and *Int-expr2*, including the values of *Int-expr1* and *Int-expr2* themselves:

```
Sequence{ 1..(6 + 4) }
Sequence{ 1..10 }
-- are both identical to
Sequence{ 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 }
```

The complete list of Collection operations is described in Clause 11 ("The OCL Standard Library").

Collections can be specified by a literal, as described above. The only other way to get a collection is by navigation. To be more precise, the only way to get a Set, OrderedSet, Sequence, or Bag is:

1. a literal, this will result in a Set, OrderedSet, Sequence, or Bag:

```
Set      {2, 4, 1, 5, 7, 13, 11, 17}
OrderedSet {1, 2, 3, 5, 7, 11, 13, 17}
Sequence {1, 2, 3, 5, 7, 11, 13, 17}
Bag      {1, 2, 3, 2, 1}
```

2. a navigation starting from a single object can result in a collection:

```
context Company inv:
    self.employee
```

3. operations on collections may result in new collections:

```
collection1->union(collection2)
```

7.6.12 Collections of Collections

OCL allows elements of collections to be collections themselves. The OCL Standard Library includes specific flattened operations for collections. These can be used to flatten collections of collections explicitly.

7.6.13 Collection Type Hierarchy and Type Conformance Rules

In addition to the type conformance rules in 7.5.5, 'Type Conformance' the following rules hold for all types, including the collection types:

- The types Set (X), Bag (X), and Sequence (X) are all subtypes of Collection (X).

Type conformance rules are as follows for the collection types:

- *Type1* conforms to *Type2* when they are identical (standard rule for all types).
- *Type1* conforms to *Type2* when it is a subtype of *Type2* (standard rule for all types).
- *Collection(Type1)* conforms to *Collection(Type2)*, when *Type1* conforms to *Type2*. This is also true for *Set(Type1)/Set(Type2)*, *Sequence(Type1)/Sequence(Type2)*, *Bag(Type1)/Bag(Type2)*.
- Type conformance is transitive: if *Type1* conforms to *Type2*, and *Type2* conforms to *Type3*, then *Type1* conforms to *Type3* (standard rule for all types).

For example, if *Bicycle* and *Car* are two separate subtypes of *Transport*:

```
Set(Bicycle) conforms to Set(Transport)
Set(Bicycle) conforms to Collection(Bicycle)
Set(Bicycle) conforms to Collection(Transport)
```

Note that Set(Bicycle) does not conform to Bag(Bicycle), nor the other way around. They are both subtypes of Collection(Bicycle) at the same level in the hierarchy.

7.6.14 Previous Values in Postconditions

As stated in 7.4.4, 'Pre- and Postconditions' OCL can be used to specify pre- and postconditions on operations and behaviors in UML. In a postcondition, the expression can refer to values of any feature of an object at two moments in time:

- the value of a feature at the start of the operation or behavior
- the value of a feature upon completion of the operation or behavior

The value of an operation call or a property navigation in a postcondition is the value upon completion of the operation. To refer to the value of a feature at the start of the operation, one has to postfix the property name with the keyword '@pre':

```
context Person::birthdayHappens()
  post: age = age@pre + 1
```

The property *age* refers to the property of the instance of Person that executes the operation. The property *age@pre* refers to the value of the property *age* of the Person that executes the operation, at the start of the operation.

In the case of an operation call, the '@pre' is postfixed to the operation name, before the parameters.

```
context Company::hireEmployee(p : Person)
  post: employees = employees@pre->including(p) and
       stockprice() = stockprice@pre() + 10
```

When the pre-value of a feature evaluates to an object, all further properties that are accessed of this object are the new values (upon completion of the operation) of this object. So:

```
a.b@pre.c -- takes the old value of property b of a, say x
           -- and then the new value of c of x.
a.b@pre.c@pre-- takes the old value of property b of a, say x
                -- and then the old value of c of x.
```

The '@pre' postfix is allowed only in OCL expressions that are part of a Postcondition, and only on invocations of the features of model classifiers. Asking for a current property of an object that has been destroyed during execution of the operation results in null. Also, referring to the previous value of an object that has been created during execution of the operation results in null.

7.6.15 Tuples

It is possible to compose several values into a *tuple*. A tuple consists of named parts, each of which can have a distinct type. Some examples of tuples are:

```
Tuple {name: String = 'John,' age: Integer = 10}
Tuple {a: Collection(Integer) = Set{1, 3, 4}, b: String = 'foo,' c: String = 'bar'}
```

This is also the way to write tuple literals in OCL; they are enclosed in curly brackets, and the parts are separated by commas. The type names are optional, and the order of the parts is unimportant. Thus:

```
Tuple {name: String = 'John,' age: Integer = 10} is equivalent to
Tuple {name = 'John,' age = 10} and to
Tuple {age = 10, name = 'John'}
```

Also, note that the values of the parts may be given by arbitrary OCL expressions, so for example we may write:

```
context Person def:
  attr statistics : Set(TupleType(company: Company, numEmployees: Integer,
                                wellpaidEmployees: Set(Person), totalSalary: Integer)) =
    managedCompanies->collect(c |
      Tuple { company: Company = c,
              numEmployees: Integer = c.employee->size(),
              wellpaidEmployees: Set(Person) = c.Job->select(salary>10000).employee->asSet(),
```

```

        totalSalary: Integer = c.Job.salary->sum()
    }
)

```

This results in a bag of tuples summarizing the company, number of employees, the best paid employees, and total salary costs of each company a person manages.

The parts of a tuple are accessed by their names, using the same dot notation that is used for accessing attributes. Thus:

```
Tuple {x: Integer = 5, y: String = 'hi'}.x = 5
```

is a true, if somewhat pointless, expression. Using the definition of statistics above, we can write:

```

context Person inv:
    statistics->sortedBy(totalSalary)->last().wellpaidEmployees->includes(self)

```

This asserts that a person is one of the best-paid employees of the company with the highest total salary that he manages. In this expression, both 'totalSalary' and 'wellpaidEmployees' are accessing tuple parts.

7.7 Collection Operations

OCIL defines many operations on the collection types. These operations are specifically meant to enable a flexible and powerful way of projecting new collections from existing ones. The different constructs are described in the following sub clauses.

7.7.1 Select and Reject Operations

Sometimes an expression using operations and navigations results in a collection, while we are interested only in a special subset of the collection. OCIL has special constructs to specify a selection from a specific collection. These are the *select* and *reject* operations. The select specifies a subset of a collection. A select is an operation on a collection and is specified using the arrow-syntax:

```
collection->select( ... )
```

The parameter of select has a special syntax that enables one to specify which elements of the collection we want to select. There are three different forms, of which the simplest one is:

```
collection->select( boolean-expression )
```

This results in a collection that contains all the elements from *collection* for which the *boolean-expression* evaluates to true. To find the result of this expression, for each element in *collection* the expression *boolean-expression* is evaluated. If this evaluates to true, the element is included in the result collection, otherwise not. As an example, the following OCIL expression specifies that the collection of all the employees older than 50 years is not empty:

```

context Company inv:
    self.employee->select(age > 50)->notEmpty()

```

The *self.employee* is of type Set(Person). The *select* takes each person from *self.employee* and evaluates *age > 50* for this person. If this results in *true*, then the person is in the result Set.

As shown in the previous example, the context for the expression in the select argument is the element of the collection on which the select is invoked. Thus the *age* property is taken in the context of a person.

In the above example, it is impossible to refer explicitly to the persons themselves; you can only refer to properties of them. To enable to refer to the persons themselves, there is a more general syntax for the select expression:

```
collection->select( v | boolean-expression-with-v )
```

The variable v is called the iterator. When the select is evaluated, v iterates over the *collection* and the *boolean-expression-with- v* is evaluated for each v . The v is a reference to the object from the collection and can be used to refer to the objects themselves from the *collection*. The two examples below are identical:

```
context Company inv:
  self.employee->select(age > 50)->notEmpty()

context Company inv:
  self.employee->select(p | p.age > 50)->notEmpty()
```

The result of the complete select is the collection of persons p for which the $p.age > 50$ evaluates to True. This amounts to a subset of *self.employee*.

As a final extension to the select syntax, the expected type of the variable v can be given. The select now is written as:

```
collection->select( v : Type | boolean-expression-with-v )
```

The meaning of this is that the objects in *collection* must be of type *Type*. The next example is identical to the previous examples:

```
context Company inv:
  self.employee.select(p : Person | p.age > 50)->notEmpty()
```

The complete select syntax now looks like one of:

```
collection->select( v : Type | boolean-expression-with-v )
collection->select( v | boolean-expression-with-v )
collection->select( boolean-expression )
```

The *reject* operation is identical to the select operation, but with reject we get the subset of all the elements of the collection for which the expression evaluates to False. The reject syntax is identical to the select syntax:

```
collection->reject( v : Type | boolean-expression-with-v )
collection->reject( v | boolean-expression-with-v )
collection->reject( boolean-expression )
```

As an example, specify that the collection of all the employees who are **not** married is empty:

```
context Company inv:
  self.employee->reject( isMarried )->isEmpty()
```

The reject operation is available in OCL for convenience, because each reject can be restated as a select with the negated expression. Therefore, the following two expressions are identical:

```
collection->reject( v : Type | boolean-expression-with-v )
collection->select( v : Type | not (boolean-expression-with-v) )
```

7.7.2 Collect Operation

As shown in the previous sub clause, the select and reject operations always result in a sub-collection of the original collection. When we want to specify a collection that is derived from some other collection, but which contains different objects from the original collection (i.e., it is not a sub-collection), we can use a *collect* operation. The collect operation uses the same syntax as the select and reject and is written as one of:

```
collection->collect( v : Type | expression-with-v )
collection->collect( v | expression-with-v )
collection->collect( expression )
```

The value of the reject operation is the collection of the results of all the evaluations of *expression-with- v* .

An example: specify the collection of *birthDates* for all employees in the context of a company. This can be written in the context of a Company object as one of:

```
self.employee->collect( birthDate )
self.employee->collect( person | person.birthDate )
self.employee->collect( person : Person | person.birthDate )
```

An important issue here is that when the source collection is a Set the resulting collection is not a Set but a Bag. Moreover, if the source collection is a Sequence or an OrderedSet, the resulting collection is a Sequence. When more than one employee has the same value for *birthDate*, this value will be an element of the resulting Bag more than once. The Bag resulting from the *collect* operation always has the same size as the original collection.

It is possible to make a Set from the Bag, by using the *asSet* property on the Bag. The following expression results in the Set of different *birthDates* from all employees of a Company:

```
self.employee->collect( birthDate )->asSet()
```

Shorthand for Collect

Because navigation through many objects is very common, there is a shorthand notation for the collect that makes the OCL expressions more readable. Instead of

```
self.employee->collect(birthdate)
```

we can also write:

```
self.employee.birthdate
```

In general, when we apply a property to a collection of Objects, then it will automatically be interpreted as a *collect* over the members of the collection with the specified property.

For any *propertyname* that is defined as a property on the objects in a collection, the following two expressions are identical:

```
collection.propertyname
collection->collect(propertyname)
```

and so are these if the property is parameterized:

```
collection.propertyname (par1, par2, ...)
collection->collect (propertyname(par1, par2, ...))
```

7.7.3 ForAll Operation

Many times a constraint is needed on all elements of a collection. The *forAll* operation in OCL allows specifying a Boolean expression, which must hold for all objects in a collection:

```
collection->forAll( v : Type | boolean-expression-with-v )
collection->forAll( v | boolean-expression-with-v )
collection->forAll( boolean-expression )
```

This *forAll* expression results in a Boolean. The result is true if the *boolean-expression-with-v* is true for all elements of *collection*. If the *boolean-expression-with-v* is false for one or more *v* in *collection*, then the complete expression evaluates to false. For example, in the context of a company:

```
context Company
inv: self.employee->forAll( age <= 65 )
inv: self.employee->forAll( p | p.age <= 65 )
inv: self.employee->forAll( p : Person | p.age <= 65 )
```

These invariants evaluate to true if the *age* property of each employee is less or equal to 65.

The *forAll* operation has an extended variant in which more than one iterator is used. Both iterators will iterate over the complete collection. Effectively this is a *forAll* on the Cartesian product of the collection with itself.

```
context Company inv:
  self.employee->forAll( e1, e2 : Person |
    e1 <> e2 implies e1.forename <> e2.forename)
```

This expression evaluates to true if the forenames of all employees are different. It is semantically equivalent to:

```
context Company inv:
  self.employee->forAll( e1 | self.employee->forAll( e2 |
    e1 <> e2 implies e1.forename <> e2.forename))
```

7.7.4 Exists Operation

Many times one needs to know whether there is at least one element in a collection for which a constraint holds. The *exists* operation in OCL allows you to specify a Boolean expression that must hold for at least one object in a collection:

```
collection->exists( v : Type | boolean-expression-with-v )
collection->exists( v | boolean-expression-with-v )
collection->exists( boolean-expression )
```

This exists operation results in a Boolean. The result is true if the *boolean-expression-with-v* is true for at least one element of *collection*. If the *boolean-expression-with-v* is false for all *v* in *collection*, then the complete expression evaluates to false. For example, in the context of a company:

```
context Company inv:
  self.employee->exists( forename = 'Jack' )

context Company inv:
  self.employee->exists( p | p.forename = 'Jack' )

context Company inv:
  self.employee->exists( p : Person | p.forename = 'Jack' )
```

These expressions evaluate to true if the *forename* property of at least one employee is equal to 'Jack.'

Similarly to *forAll* expression an exists expression may declare multiple iterators.

7.7.5 Closure Operation

The iterators described in the preceding sections return results from the elements of a collection. The *closure* supports returning results from the elements of a collection, the elements of the elements of a collection, the elements of the elements of the elements of a collection, and so forth. This can be useful for iterating over a transitive relationship such as a UML generalization. *closure* operation uses the same syntax as the *select* and *reject* iterators and is written as one of

```
source>closure( v : Type | expression-with-v )
source>closure( v | expression-with-v )
source>closure( expression )
```

The returned collection of the *closure* iteration is an accumulation of the source, and the collections resulting from the recursive invocation of *expression-with-v* in which *v* is associated exactly once with each distinct element of the returned collection. The iteration terminates when *expression-with-v* returns empty collections or collections containing only already accumulated elements. The collection type of the *result* collection is the unique form (Set or OrderedSet) of the original *source* collection. If the *source* collection is ordered, the *result* is in depth first preorder. The *result* satisfies the postcondition:

```
post: let sourceAndResult : Set(Type) = source->asSet()->union(result) in
sourceAndResult = sourceAndResult->collect(expression)
```

For a simple parent-children relationship and known parents

```
parents->closure(children)
```

computes the set of parents.children, parents.children.children, parents.children.children.children etc.

In the opposite direction

```
self->asOrderedSet()->closure(mother)
```

computes the maternal line.

For a more complex relationship such as UML Classifier generalization

```
aClassifier.generalization()->closure(general.generalization).general()->including(aClassifier)
```

computes the set comprising aClassifier and all its generalizations. The closure recurses over the Generalizations to compute the transitive set of all Generalizations. The generalized classifier is collected from each of these before including the originating aClassifier in the result.

As with all other iterators, self remains unchanged throughout the recursion, and an implicit source attempts to resolve features against iterators.

7.7.6 Iterate Operation

The *iterate* operation is slightly more complicated, but is very generic. The operations *reject*, *select*, *forAll*, *exists*, *collect* can all be described in terms of *iterate*. An accumulation builds one value by iterating over a collection.

```
collection->iterate( elem : Type; acc : Type = <expression> |
expression-with-elem-and-acc )
```

The variable *elem* is the iterator, as in the definition of *select*, *forAll*, etc. The variable *acc* is the accumulator. The accumulator gets an initial value *<expression>*. When the *iterate* is evaluated, *elem* iterates over the *collection* and the *expression-with-elem-and-acc* is evaluated for each *elem*. After each evaluation of *expression-with-elem-and-acc*, its value is assigned to *acc*. In this way, the value of *acc* is built up during the iteration of the collection. The collect operation described in terms of *iterate* will look like:

```
collection->collect(x : T | x.property)
-- is identical to:
collection->iterate(x : T; acc : T2 = Bag{} |
acc->including(x.property))
```

Or written in Java-like pseudocode the result of the *iterate* can be calculated as:

```
iterate(elem : T; acc : T2 = value)
{
  acc = value;
  for(Enumeration e = collection.elements() ; e.hasMoreElements(); ){
    elem = e.nextElement();
    acc.add(<expression-with-elem-and-acc>)
  }
  return acc;
}
```

Although the Java pseudo code uses a ‘next element,’ the *iterate* operation is defined not only for Sequence, but for each collection type. The order of the iteration through the elements in the collection is not defined for Set and Bag. For a Sequence the order is the order of the elements in the sequence.

7.8 Messages in OCL

This sub clause contains some examples of the concrete syntax and explains the finer details of the message expression. In earlier versions the phrase “actions in OCL” was used, but message was found to capture the meaning more precisely.

7.8.1 Calling operations and sending signals

To specify that communication has taken place, the `hasSent` (‘^’) operator is used:

```
context Subject::hasChanged()
post: observer^update(12, 14)
```

The *observer^update(12, 14)* results in true if an update message with arguments 12 and 14 was sent to *observer* during the execution of the operation. *Update()* is either an *Operation* that is defined in the class of *observer*, or it is a *Signal* specified in the UML model. The argument(s) of the message expression (12 and 14 in this example) must conform to the parameters of the operation/signal definition.

If the actual arguments of the operation/signal are not known, or not restricted in any way, it can be left unspecified. This is shown by using a question mark. Following the question mark is an optional type, which may be needed to find the correct operation when the same operation exists with different parameter types.

```
context Subject::hasChanged()
post: observer^update(? : Integer, ? : Integer)
```

This example states that the message update has been sent to *observer*, but that the values of the parameters are not known.

OCL also defines a special *OclMessage* type. One can get the actual OclMessages through the message operator: ^^.

```
context Subject::hasChanged()
post: observer^^update(12, 14)
```

This results in the Sequence of messages sent. Each element of the collection is an instance of *OclMessage*. In the remainder of the constraint one can refer to the parameters of the operation using their formal parameter name from the operation definition. If the operation update has been defined with formal parameters named *i* and *j*, then we can write:

```
context Subject::hasChanged()
post: let messages : Sequence(OclMessage) = observer^^update(? : Integer, ? : Integer) in
    messages->notEmpty() and
    messages->exists( m | m.i > 0 and m.j >= m.i )
```

The value of the parameter *i* is not known, but it must be greater than zero and the value of parameter *j* must be larger or equal to *i*.

Because the ^^ operator results in an instance of *OclMessage*, the message expression can also be used to specify collections of messages sent to different targets. For an observer pattern we can write:

```
context Subject::hasChanged()
post: let messages : Sequence(OclMessage) =
    observers->collect(o | o^^update(? : Integer, ? : Integer) ) in
    messages->forAll(m | m.i <= m.j )
```

Messages is now a set of *OclMessage* instances, where every *OclMessage* instance has one of the *observers* as a *target*.

7.8.2 Accessing result values

A signal sent message is by definition asynchronous, so there never is a return value. If there is a logical return value it must be modeled as a separate signal message. Yet, for an operation call there is a potential return value. This is only available if the operation has already returned (not necessary if the operation call is asynchronous), and it specifies a return type in its definition. The standard operation *result()* of *OclMessage* contains the return value of the called operation. If *getMoney(...)* is an operation on *Company* that returns a boolean, as in *Company::getMoney(amount : Integer) : Boolean*, we can write:

```
context Person::giveSalary(amount : Integer)
post: let message : OclMessage = company^getMoney(amount) in
    message.hasReturned()           -- getMoney was sent and returned
    and
    message.result() = true          -- the getMoney call returned true
```

As with the previous example we can also access a collection of return values from a collection of *OclMessages*. If *message.hasReturned()* is false, then *message.result()* will be *invalid*.

7.8.3 An example

This sub clause shows an example of using the OCL message expression.

The Example and Problem

Suppose we have built a component, which takes any form of input and transforms it into garbage (aka encrypts it). The component *GarbageCan* uses an interface *UsefullInformationProvider* that must be implemented by users of the component to provide the input. The operation *getNextPieceOfGarbage* of *GarbageCan* can then be used to retrieve the garbled data. Figure 7.6 shows the component's class diagram. Note that none of the operations are marked as queries.




Figure 7.6 - OclMessageExample

When selling the component, we do not want to give the source code to our customers. However, we want to specify the component's behavior as precisely as possible. So, for example, we want to specify, what *getNextPieceOfGarbage* does. Note that we cannot write:

```
context GarbageCan::getNextPieceOfGarbage() : Integer
post: result = (datasource.getNextPieceOfData() * .7683425 + 10000) / 20 + 3
```

because *UsefulInformationProvider::getNextPieceOfData()* is not a query (e.g., it may increase some internal pointer so that it can return the next piece of data at the next call). Still we would like to say something about how the garbage is derived from the original data.

The solution

To solve this problem, we can use an *OclMessage* to represent the call to *getNextPieceOfData*. This allows us to check for the result. Note that we need to demand that the call has returned before accessing the result:

```
context GarbageCan::getNextPieceOfGarbage() : Integer
post: let message : OclMessage = datasource^^getNextPieceOfData()->first() in
    message.hasReturned()
and
    result = (message.result() * .7683425 + 10000) / 20 + 3
```

7.9 Resolving Properties

For any property (attribute, operation, or navigation) the full notation includes the object of which the property is taken. As seen in 7.4.3, 'Invariants' *self* can be left implicit, and so can the iterator variables in collection operations. At any place in an expression, when an iterator is left out, an implicit iterator-variable is introduced. For example in:

```
context Person inv:
    employer->forAll( employee->exists( lastName = name) )
```

three implicit variables are introduced. The first is *self*, which is always the instance from which the constraint starts. Secondly an implicit iterator is introduced by the *forAll* and third by the *exists*. The implicit iterator variables are unnamed. The properties *employer*, *employee*, *lastName*, and *name* all have the object on which they are applied left out. Resolving these goes as follows:

- at the place of *employer* there is one implicit variable: *self : Person*. Therefore *employer* must be a property of *self*.
- at the place of *employee* there are two implicit variables: *self : Person* and *iter1 : Company*. Therefore *employee* must be a property of either *self* or *iter1*. If *employee* is a property of both *self* and *iter1*, then it is defined to belong to the variable in the most inner scope, which is *iter1*.
- at the place of *lastName* and *name* there are three implicit variables: *self : Person*, *iter1 : Company* and *iter2 : Person*. Therefore *lastName* and *name* must both be a property of either *self* or *iter1* or *iter2*. In the UML model property *name* is a property of *iter1*. However, *lastName* is a property of both *self* and *iter2*. This is ambiguous and therefore the *lastName* refers to the variable in the most inner scope, which is *iter2*.

Both of the following invariant constraints are correct, but have a different meaning:

```
context Person
inv: employer->forAll( employee->exists( p | p.lastName = name) )
inv: employer->forAll( employee->exists( self.lastName = name) )
```

A closure iteration may introduce an implicit iterator-variable at each level of recursion and so multiple iterator-variable candidates for consideration as the implicit *self*. Since all candidates have the same static type, it is only the least deeply nested candidate, with respect to the iteration body, that need be considered as the implicit iterator-variable for a closure.

8 Abstract Syntax

8.1 Introduction

This clause describes the abstract syntax of the OCL. In this abstract syntax a number of metaclasses from the UML metamodel are imported. These metaclasses are shown in the models with a transparent fill color. All metaclasses defined as part of the OCL abstract syntax are shown with a light gray background.

The abstract syntax as described below defines the concepts that are part of the OCL using a MOF compliant metamodel. The abstract syntax is divided into several packages.

- The *Types* package describes the concepts that define the type system of OCL. It shows which types are predefined in OCL and which types are deduced from the UML models.
- The *Expressions* package describes the structure of OCL expressions.

8.2 The Types Package

OCL is a typed language. Each expression has a type that is either explicitly declared or can be statically derived. Evaluation of the expression yields a value of this type. Therefore, before we can define expressions, we have to provide a model for the concept of type. A metamodel for OCL types is shown in this sub clause. Note that instances of the classes in the metamodel are the types themselves (e.g., Integer) not instances of the domain they represent (e.g., -15, 0, 2, 3).

The model in Figure 8.1 shows the OCL types. The basic type is the UML Classifier, which includes all subtypes of Classifier from the UML Superstructure.

In the model, the *CollectionType* (and its subclasses) and the *TupleType* are special. One can never instantiate all collection types, because there is an infinite number, especially when nested collections are taken into account. Conceptually all these types do exist, but such a type should be (lazily) instantiated by a tool, whenever it is needed in an expression. For convenience an instance representing a collection type or a tuple type may be replicated in different namespaces (such as in a top-level package or within the expression referencing it), however they represent semantically the same type.




Figure 8.1 - Abstract Syntax Kernel Metamodel for OCL Types

AnyType

AnyType is the metaclass of the special type *OclAny*, which is the type to which all other types conform. *OclAny* is the sole instance of *AnyType*. This metaclass allows defining the special property of being the generalization of all other Classifiers, including Classes, DataTypes, and PrimitiveTypes.

BagType

BagType is a collection type that describes a multiset of elements where each element may occur multiple times in the bag. The elements are unordered. Part of a *BagType* is the declaration of the type of its elements.

CollectionType

CollectionType describes a list of elements of a particular given type. *CollectionType* is a concrete metaclass whose instances are the family of abstract *Collection(T)* data types. Its subclasses are *SetType*, *OrderedSetType*, *SequenceType*, and *BagType*, whose instances are the concrete *Set(T)*, *OrderedSet(T)*, *Sequence(T)*, and *Bag(T)*, data types, respectively.

Part of every collection type is the declaration of the type of its elements (i.e., a collection type is parameterized with an element type). In the metamodel, this is shown as an association from *CollectionType* to *Classifier*. Note that there is no restriction on the element type of a collection type. This means in particular that a collection type may be parameterized with other collection types allowing collections to be nested arbitrarily deep.

Associations

elementType	The type of the elements in a collection. All elements in a collection must conform to this type.
-------------	---

InvalidType

InvalidType represents a type that conforms to all types except the *VoidType* type. The only instance of *InvalidType* is *Invalid*, which is further defined in the standard library. Furthermore *Invalid* has exactly one runtime instance identified as *OclInvalid*.

MessageType

MessageType describes ocl messages. Similar to the collection types, *MessageType* describes a set of types in the standard library. Part of every *MessageType* is a reference to the declaration of the type of its operation or signal, i.e., an ocl message type is *parameterized* with an operation or signal. In the metamodel, this is shown as an association from *MessageType* to *Operation* and to *Signal*. *MessageType* is part of the abstract syntax of OCL, residing on M2 level. Its instances, called *OclMessage*, and subtypes of *OclMessage*, reside on M1 level.

Associations

referredSignal	The <i>Signal</i> that is sent by the message.
referredOperation	The <i>Operation</i> that is called by the message.

OrderedSetType

OrderedSetType is a collection type that describes a set of elements where each distinct element occurs only once in the set. The elements are ordered by their position in the sequence. Part of an *OrderedSetType* is the declaration of the type of its elements.

SequenceType

SequenceType is a collection type that describes a list of elements where each element may occur multiple times in the sequence. The elements are ordered by their position in the sequence. Part of a *SequenceType* is the declaration of the type of its elements.

SetType

SetType is a collection type that describes a set of elements where each distinct element occurs only once in the set. The elements are not ordered. Part of a *SetType* is the declaration of the type of its elements.

TemplateParameterType

A *TemplateParameterType* is used to refer to generic types in parameterized definitions. It is used in the standard library to represent the parameterized collection operations. A *TemplateParameterType* is usually named “T” (or “T2,” “T3,” and so on, when more than one type parameter is involved).

The *TemplateParameterType* is a sub-class of *Classifier*.

Attributes

specification An un-interpreted opaque definition of the template parameter type.

TupleType

TupleType (informally known as record type or struct) combines different types into a single aggregate type. The parts of a *TupleType* are described by its attributes, each having a name and a type. There is no restriction on the kind of types that can be used as part of a tuple. In particular, a *TupleType* may contain other tuple types and collection types. Each attribute of a *TupleType* represents a single feature of a *TupleType*. Each part is uniquely identified by its name.

VoidType

VoidType is the metaclass of the *OclVoid* type that conforms to all types except the *OclInvalid* type. The only instance of *VoidType* is *OclVoid*, which is further defined in the standard library. Furthermore *OclVoid* has exactly one instance called *null* - corresponding to the UML NullLiteral literal specification - and representing the *absence of value*. Note that in contrast with *invalid*, *null* is a valid value and as such can be owned by collections.

8.2.1 Type Conformance

The type conformance rules are formally underpinned in the Semantics sub clause of the specification. To ensure that the rules are accessible to UML modelers they are specified in this sub clause using OCL. For this, the additional operation *conformsTo*(*c* : *Classifier*) : *Boolean* is defined on *Classifier*. It evaluates to true, if the self *Classifier* conforms to the argument *c*. The following OCL statements define type conformance for individual types.

BagType

[3] [1] Different bag types conform to each other if their element types conform to each other.

```
context BagType
inv: BagType.allInstances()->forAll(b |
    self.elementType.conformsTo(b.elementType) implies self.conformsTo(b))
```

Classifier

[4] [1] Conformance is a transitive relationship.

```
context Classifier
inv Transitivity: Classifier.allInstances()->forAll(x|Classifier.allInstances()
    ->forAll(y|
        (self.conformsTo(x) and x.conformsTo(y)) implies self.conformsTo(y)))
```

[5] [2] Classes conform to superclasses and interfaces that they realize.

```
context Class
inv : self.generalization.general->forAll (p |
    (p.oclIsKindOf(Class) or p.oclIsKindOf(Interface)) implies
        self.conformsTo(p.oclAsType(Classifier)))
```

[6] [3] Interfaces conforms to super interfaces.

```
context Interface
inv : self.generalization.general->forAll (p |
    p.oclIsKindOf(Interface) implies self.conformsTo(p.oclAsType(Interface)))
```

[7] [4] The Conforms operation between Types is reflexive, a Classifier always conform to itself.

```
context Classifier
inv: self.conformsTo(self)
```

[8] [5] The Conforms operation between Types is anti-symmetric.

```
context Classifier
inv: Classifier.allInstances()->forAll(t1, t2 |
    (t1.conformsTo(t2) and t2.conformsTo(t1)) implies t1 = t2)
```

CollectionType

[9] [1] Specific collection types conform to collection type.

```
context CollectionType
inv: -- all instances of SetType, SequenceType, BagType conform to a
    -- CollectionType if the elementTypes conform
    CollectionType.allInstances()->forAll (c |
        c.oclIsTypeOf(CollectionType) and
        self.elementType.conformsTo(c.elementType) implies
        self.conformsTo(c))
```

[10] [2] Collections do not conform to any primitive type.

```
context CollectionType
inv: PrimitiveType.allInstances()->forAll (p | not self.conformsTo(p))
```

[11] [3] Collections of non-conforming types do not conform.

```
context CollectionType
inv: CollectionType.allInstances()->forAll (c |
    (not self.elementType.conformsTo (c.elementType)) implies (not self.conformsTo (c)))
```

InvalidType

[12] [1] OclInvalid conforms to all other types.

```
context InvalidType
inv: Classifier.allInstances()->forAll (c | self.conformsTo (c))
```

OrderedSetType

[13] [1] Different ordered set types conform to each other if their element types conform to each other.

```
context OrderedSetType
inv: OrderedSetType.allInstances()->forAll(s |
    self.elementType.conformsTo(s.elementType) implies self.conformsTo(s))
```

PrimitiveType

[14] [1] Integer conforms to Real.

```
context PrimitiveType
inv: (self.name = 'Integer') implies
    PrimitiveType.allInstances()->forAll (p | (p.name = 'Real') implies
        (self.conformsTo(p)))
```

[15] [2] UnlimitedNatural conforms to Integer.

```
context PrimitiveType
inv: (self.name = 'UnlimitedNatural') implies
    PrimitiveType.allInstances()->forAll (p | (p.name = 'Integer') implies
        (self.conformsTo(p)))
```

Note that * is an invalid Integer and so conversion of * to Integer yields invalid whose type conforms to all types.

SequenceType

[16][1] Different sequence types conform to each other if their element types conform to each other.

```
context SequenceType
inv: SequenceType.allInstances()->forAll(s |
    self.elementType.conformsTo(s.elementType) implies self.conformsTo(s))
```

SetType

[17][1] Different set types conform to each other if their element types conform to each other.

```
context SetType
inv: SetType.allInstances()->forAll(s |
    self.elementType.conformsTo(s.elementType) implies self.conformsTo(s))
```

TupleType

[18][1] Tuple types conform to each other when their names and types conform to each other. Note that *allProperties* is an additional operation in the UML.

```
context TupleType
inv: TupleType.allInstances()->forAll (t |
    ( t.allProperties()->forAll (tp |
        -- make sure at least one tuplepart has the same name
        -- (uniqueness of tuplepart names will ensure that not two
        -- tupleparts have the same name within one tuple)
        self.allProperties()->exists(stp|stp.name = tp.name) and
        -- make sure that all tupleparts with the same name conforms.
        self.allProperties()->forAll(stp | (stp.name = tp.name) implies
            stp.type.conformsTo(tp.type))
    )
    implies
        self.conformsTo(t)
)
```

VoidType

[19][1] OclVoid conforms to all other types except OclInvalid.

```
context VoidType
inv: Classifier.allInstances()->forAll (c | not c.oclIsTypeOf(OclInvalid) implies self.conformsTo (c))
```

8.2.2 Operations and Well-formedness Rules for the Types Package

BagType

[20][1] The name of a bag type is “Bag” followed by the element type’s name in parentheses.

```
context BagType
inv: self.name = 'Bag(' + self.elementType.name + ')'
```

BooleanType

allInstances() : Set(Boolean)

Returns Set{true,false}.