
Meta Object Facility (MOF) Specification

Version 1.3
New Edition: March 2000

Copyright 1997-1999, DSTC (Cooperative Research Centre for Enterprise Distributed Systems Technology)
Copyright 1997-1999, Electronic Data Systems
Copyright 1997-1999, IBM Corporation
Copyright 1997-1999, International Computers Limited
Copyright 1997-1999, Objectivity Inc.
Copyright 2000, Object Management Group
Copyright 1997-1999, Oracle Corporation
Copyright 1997-1999, Platinum Technology Inc.
Copyright 1997-1999, Rational Software Corporation
Copyright 1997-1999, System Software Associates
Copyright 1997-1999, Unisys Corporation

The companies listed above have granted to the Object Management Group, Inc. (OMG) a nonexclusive, royalty-free, paid up, worldwide license to copy and distribute this document and to modify this document and distribute copies of the modified version. Each of the copyright holders listed above has agreed that no person shall be deemed to have infringed the copyright in the included material of any such copyright holder by reason of having used the specification set forth herein or having conformed any computer software to the specification.

PATENT

The attention of adopters is directed to the possibility that compliance with or adoption of OMG specifications may require use of an invention covered by patent rights. OMG shall not be responsible for identifying patents for which a license may be required by any OMG specification, or for conducting legal inquiries into the legal validity or scope of those patents that are brought to its attention. OMG specifications are prospective and advisory only. Prospective users are responsible for protecting themselves against liability for infringement of patents.

NOTICE

The information contained in this document is subject to change without notice. The material in this document details an Object Management Group specification in accordance with the license and notices set forth on this page. This document does not represent a commitment to implement any portion of this specification in any company's products.

WHILE THE INFORMATION IN THIS PUBLICATION IS BELIEVED TO BE ACCURATE, THE OBJECT MANAGEMENT GROUP AND THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS MATERIAL INCLUDING, BUT NOT LIMITED TO ANY WARRANTY OF TITLE OR OWNERSHIP, IMPLIED WARRANTY OF MERCHANTABILITY OR WARRANTY OF FITNESS FOR PARTICULAR PURPOSE OR USE. In no event shall The Object Management Group or any of the companies listed above be liable for errors contained herein or for indirect, incidental, special, consequential, reliance or cover damages, including loss of profits, revenue, data or use, incurred by any user or any third party. The copyright holders listed above acknowledge that the Object Management Group (acting itself or through its designees) is and shall at all times be the sole entity that may authorize developers, suppliers and sellers of computer software to use certification marks, trademarks or other special designations to indicate compliance with these materials. This document contains information which is protected by copyright. All Rights Reserved. No part of this work covered by copyright herein may be reproduced or used in any form or by any means--graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems--without permission of the copyright owner.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Right in Technical Data and Computer Software Clause at DFARS 252.227.7013 OMG[®] and Object Management are registered trademarks of the Object Management Group, Inc. Object Request Broker, OMG IDL, ORB, CORBA, CORBAfacilities, CORBAservices, COSS, and IIOP are trademarks of the Object Management Group, Inc. X/Open is a trademark of X/Open Company Ltd.

The UML logo is a trademark of Rational Software Corp.

ISSUE REPORTING

All OMG specifications are subject to continuous review and improvement. As part of this process we encourage readers to report any ambiguities, inconsistencies, or inaccuracies they may find by completing the issue reporting form at <http://www.omg.org/library/issuerpt.htm>.

Table of Contents

Preface	ix
1. MOF Overview	1-1
1.1 Overview	1-1
1.2 Software Development Scenarios.....	1-2
1.3 Type Management Scenarios	1-4
1.4 Information Management Scenarios.....	1-6
1.5 Data Warehouse Management Scenarios	1-6
2. MOF Conceptual Overview	2-1
2.1 Overview	2-1
2.2 Meta-data Architectures	2-1
2.2.1 Four Layer Meta-data Architectures.....	2-2
2.2.2 The MOF Meta-data Architecture	2-3
2.2.3 MOF Meta-modeling Terminology.....	2-5
2.3 The MOF Model - Meta-modeling Constructs	2-5
2.3.1 Classes	2-6
2.3.2 Associations	2-9
2.3.3 Aggregation	2-11
2.3.4 References	2-12
2.3.5 DataTypes.....	2-14
2.3.6 Packages	2-14
2.3.7 Constraints and Consistency.....	2-17
2.3.8 Miscellaneous Meta-modeling Constructs	2-20
2.4 Meta-models and Mappings	2-21
2.4.1 Abstract and Concrete Mappings	2-21

2.4.2	CORBA Meta-data Services - The MOF IDL Mapping	2-22
2.4.3	Meta-data Interchange - The MOF XML Mapping	2-22
2.4.4	Mappings of the MOF Model	2-22
3.	MOF Model and Interfaces	3-1
3.1	Overview	3-1
3.2	How the MOF Model is Described.	3-2
3.2.1	Classes	3-3
3.2.2	Associations	3-8
3.2.3	DataTypes	3-9
3.2.4	Exceptions	3-10
3.2.5	Constants	3-10
3.2.6	Constraints	3-10
3.2.7	UML Diagrams.	3-10
3.3	The Structure of the MOF Model.	3-11
3.3.1	The MOF Model Package.	3-11
3.3.2	The MOF Model Service IDL	3-12
3.3.3	The MOF Model Structure	3-12
3.3.4	The MOF Model Containment Hierarchy.	3-14
3.4	MOF Model Classes.	3-15
3.4.1	ModelElement Class.	3-15
3.4.2	Namespace Class	3-21
3.4.3	GeneralizableElement Class	3-24
3.4.4	TypedElement.	3-29
3.4.5	Classifier.	3-31
3.4.6	Class	3-31
3.4.7	DataType	3-33
3.4.8	TypeAlias	3-35
3.4.9	Feature Class	3-36
3.4.10	StructuralFeature	3-38
3.4.11	Attribute	3-40
3.4.12	Reference	3-41
3.4.13	BehavioralFeature Class.	3-43
3.4.14	Operation	3-44
3.4.15	Exception	3-46
3.4.16	Association	3-47
3.4.17	AssociationEnd.	3-49
3.4.18	Package.	3-53
3.4.19	Import.	3-55

	3.4.20	Parameter	3-57
	3.4.21	Constraint	3-59
	3.4.22	Constant Class	3-62
	3.4.23	Tag	3-63
3.5		MOF Model Associations	3-65
	3.5.1	Contains	3-65
	3.5.2	Generalizes	3-66
	3.5.3	RefersTo	3-68
	3.5.4	Exposes	3-69
	3.5.5	IsOfType	3-70
	3.5.6	CanRaise	3-71
	3.5.7	Aliases	3-72
	3.5.8	Constrains	3-73
	3.5.9	DependsOn	3-74
	3.5.10	AttachesTo	3-76
3.6		\MOF Model Data Types	3-77
	3.6.1	CORBA Basic Types	3-77
	3.6.2	NameType	3-77
	3.6.3	AnnotationType	3-77
	3.6.4	TypeDescriptor	3-77
	3.6.5	MultiplicityType	3-78
	3.6.6	VisibilityKind	3-79
	3.6.7	DepthKind	3-79
	3.6.8	DirectionKind	3-80
	3.6.9	ScopeKind	3-80
	3.6.10	AggregationKind	3-80
	3.6.11	EvaluationKind	3-80
	3.6.12	DependencyKind	3-81
	3.6.13	FormatType	3-81
	3.6.14	LiteralType	3-81
	3.6.15	VerifyResultKind	3-82
	3.6.16	ViolationType	3-82
3.7		MOF Model Exceptions	3-83
	3.7.1	NameNotFound	3-83
	3.7.2	NameNotResolved	3-83
	3.7.3	ObjectNotExternalizable	3-84
	3.7.4	FormatNotSupported	3-84
	3.7.5	IllformedExternalizedObject	3-84
3.8		MOF Model Constants	3-85
	3.8.1	Unbounded	3-85

3.8.2	The Standard DependencyKinds	3-85
3.9	MOF Model Constraints	3-86
3.9.1	MOF Model Constraints and other M2 Level Semantics	3-86
3.9.2	Notational Conventions	3-86
3.9.3	OCL Usage in the MOF Model specification . .	3-88
3.9.4	The MOF Model Constraints	3-91
3.9.5	Semantic specifications for some Operations, derived Attributes and Derived Associations . .	3-106
3.9.6	OCL Helper functions	3-112
4.	The MOF Abstract Mapping	4-1
4.1	Overview	4-1
4.2	MOF Values	4-2
4.2.1	Semantics of Equality for MOF Values	4-3
4.3	Semantics of Class Instances	4-3
4.4	Semantics of Attributes	4-4
4.4.1	Attribute name and type	4-4
4.4.2	Multiplicity	4-5
4.4.3	Scope	4-6
4.4.4	Is_derived	4-6
4.4.5	Aggregation	4-6
4.4.6	Visibility and is_changeable	4-7
4.5	Package Composition	4-7
4.5.1	Package Nesting	4-7
4.5.2	Package Generalization	4-8
4.5.3	Package Importation	4-8
4.5.4	Package Clustering	4-8
4.6	Extents	4-9
4.6.1	The Purpose of Extents	4-9
4.6.2	Class Extents	4-10
4.6.3	Association Extents	4-10
4.6.4	Package Extents	4-11
4.7	Semantics of Associations	4-12
4.7.1	MOF Associations in UML notation	4-14
4.7.2	Core Association Semantics	4-15
4.7.3	AssociationEnd Changeability	4-17
4.7.4	AssociationEnd Navigability	4-17
4.7.5	Association Aggregation	4-17
4.7.6	Derived Associations	4-17

4.8	Aggregation Semantics	4-17
	4.8.1 Aggregation “none”	4-18
	4.8.2 Aggregation “composite”	4-18
	4.8.3 Aggregation “shared”	4-18
4.9	Closure Rules	4-18
	4.9.1 The Reference Closure Rule	4-19
	4.9.2 The Composition Closure Rule	4-20
4.10	Recommended Copy Semantics	4-21
4.11	Computational Semantics	4-22
	4.11.1 A Style Guide for Metadata Computational Semantics	4-22
	4.11.2 Access operations should not change metadata	4-23
	4.11.3 Update operations should only change the nominated metadata	4-23
	4.11.4 Derived Elements should behave like non-derived Elements	4-23
	4.11.5 Constraint evaluation should not have side-effects	4-24
	4.11.6 Access operations should avoid raising Constraint exceptions	4-24
5.	MOF to IDL Mapping	5-1
5.1	Overview	5-1
5.2	Meta Objects and Interfaces	5-2
	5.2.1 Meta Object Type Overview	5-2
	5.2.2 The Meta Object Interface Hierarchy	5-4
5.3	Computational Semantics for the IDL Mapping	5-6
	5.3.1 Value Types and Equality in the IDL Mapping	5-6
	5.3.2 Lifecycle Semantics for the IDL Mapping	5-8
	5.3.3 Association Access and Update Semantics for the IDL Mapping	5-11
	5.3.4 Attribute Access and Update Semantics for the IDL Mapping	5-15
	5.3.5 Reference Semantics for the IDL Mapping	5-21
	5.3.6 Cluster Semantics for the IDL Mapping	5-22
	5.3.7 Atomicity Semantics for the IDL Mapping	5-22
	5.3.8 The Supertype Closure Rule	5-22
	5.3.9 Copy Semantics for the IDL Mapping	5-23
5.4	Exception Framework	5-24
	5.4.1 Error_kind string values	5-25
	5.4.2 Structural Errors	5-26

Contents

5.4.3	Constraint Errors	5-29
5.4.4	Semantic Errors	5-29
5.4.5	Usage Errors	5-30
5.4.6	Reflective Errors	5-31
5.5	Preconditions for IDL Generation	5-33
5.6	Standard Tags for the IDL Mapping	5-35
5.6.1	Tags for Specifying IDL #pragma prefix	5-36
5.6.2	Tags for Providing Substitute Identifiers	5-37
5.6.3	Tags for Specifying IDL Inheritance	5-37
5.7	Generated IDL Issues	5-39
5.7.1	Generated IDL Identifiers	5-39
5.7.2	Generation Rules for Collection Types	5-42
5.7.3	IDL Identifier Qualification	5-44
5.7.4	File Organization and #include statements	5-45
5.8	IDL Mapping Templates	5-45
5.8.1	Template Notation	5-46
5.8.2	Package Module Template	5-46
5.8.3	Package Factory Template	5-49
5.8.4	Package Template	5-51
5.8.5	Class Forward Declaration Template	5-53
5.8.6	Class Template	5-54
5.8.7	Class Proxy Template	5-54
5.8.8	Instance Template	5-56
5.8.9	Class Create Template	5-57
5.8.10	Association Template	5-59
5.8.11	Attribute Template	5-68
5.8.12	Reference Template	5-77
5.8.13	Operation Template	5-86
5.8.14	Exception Template	5-89
5.8.15	Constant Template	5-89
5.8.16	DataType Template	5-90
5.8.17	Constraint Template	5-91
5.8.18	Annotation Template	5-92
6.	The Reflective Module	6-1
6.1	Introduction	6-1
6.2	The Reflective Interfaces	6-3
6.2.1	Reflective Argument Encoding Patterns	6-3
6.2.2	Reflective::RefObject	6-9
6.2.3	Reflective::RefAssociation	6-23

6.2.4	Reflective::RefPackage.....	6-28
6.3	The CORBA IDL for the Reflective Interfaces	6-30
6.3.1	Introduction	6-30
6.3.2	Data Types.....	6-30
	Appendix A - XMI for the MOF	A-1
	Appendix B - MOF IDL Summary	B-1
	Appendix C - MODL Description of the MOF	C-1
	Appendix D - MOF Implementation Requirements	D-1
	Appendix E - Future Directions for the MOF.....	E-1

Contents

Preface

About the Object Management Group

The Object Management Group, Inc. (OMG) is an international organization supported by over 800 members, including information system vendors, software developers and users. Founded in 1989, the OMG promotes the theory and practice of object-oriented technology in software development. The organization's charter includes the establishment of industry guidelines and object management specifications to provide a common framework for application development. Primary goals are the reusability, portability, and interoperability of object-based software in distributed, heterogeneous environments. Conformance to these specifications will make it possible to develop a heterogeneous applications environment across all major hardware platforms and operating systems.

OMG's objectives are to foster the growth of object technology and influence its direction by establishing the Object Management Architecture (OMA). The OMA provides the conceptual infrastructure upon which all OMG specifications are based.

Associated OMG Documents

The CORBA documentation set includes the following:

- *Object Management Architecture Guide* defines the OMG's technical objectives and terminology and describes the conceptual models upon which OMG standards are based. It defines the umbrella architecture for the OMG standards. It also provides information about the policies and procedures of OMG, such as how standards are proposed, evaluated, and accepted.
- *CORBA: Common Object Request Broker Architecture and Specification* contains the architecture and specifications for the Object Request Broker.
- *CORBA services: Common Object Services Specification* contains specifications for OMG's Object Services.

-
- *CORBA facilities: Common Facilities Specification* is a collection of services that many applications may share, but which are not as fundamental as the Object Services. For instance, a system management or electronic mail facility could be classified as a common facility. Common Facilities are used by most systems.

The OMG collects information for each book in the documentation set by issuing Requests for Information, Requests for Proposals, and Requests for Comment and, with its membership, evaluating the responses. Specifications are adopted as standards only when representatives of the OMG membership accept them as such by vote. (The policies and procedures of the OMG are described in detail in the *Object Management Architecture Guide*.)

OMG formal documents are available from our web site in PostScript and PDF format. To obtain print-on-demand books in the documentation set or other OMG publications, contact the Object Management Group, Inc. at:

OMG Headquarters
250 First Avenue, Suite 201
Needham, MA 02494
USA
Tel: +1-781-444-0404
Fax: +1-781-444-0320
pubs@omg.org
<http://www.omg.org>

Introduction to OMG Modeling

This document describes the OMG standards for modeling distributed software architectures and systems along with their CORBA Interfaces. It is composed of two complementary specifications:

- Unified Modeling Language Specification
- Meta-Object Facility Specification

The Unified Modeling Language (UML) Specification defines a graphical language for visualizing, specifying, constructing, and documenting the artifacts of distributed object systems. The specification includes the formal definition of a common Object Analysis and Design (OA&D) metamodel, a graphic notation, and a CORBA IDL facility that supports model interchange between OA&D tools and metadata repositories. The UML provides the foundation for specifying and sharing CORBA-based distributed object models.

The Meta-Object Facility (MOF) Specification defines a set of CORBA IDL interfaces that can be used to define and manipulate a set of interoperable metamodels and their corresponding models. These interoperable metamodels include the UML metamodel, the MOF meta-metamodel, as well as future OMG adopted technologies that will be specified using metamodels. The MOF provides the infrastructure for implementing CORBA-based design and reuse repositories. The MOF specifies precise mapping

rules that enable the CORBA interfaces for metamodels to be automatically generated, thus encouraging consistency in manipulating metadata in all phases of the distributed application development cycle.

Since the UML and MOF are based on a four-layer metamodel architecture it is essential that the metamodels for each facility are architecturally aligned. For a description of the four layer metamodel architecture, please refer to Section 2.2, “Meta-data Architectures,” on page 2-1 in the MOF Specification. In order to achieve architectural alignment considerable effort has been expended so that the UML and MOF share the same core semantics. This alignment allows the MOF to reuse the UML notation for visualizing metamodels. In those areas where semantic differences are required, well-defined mapping rules are provided between the metamodels. The OMG distributed repository architecture, which integrates UML and MOF with CORBA is described in Section 0.5, “Resolution of Technical Criteria in the Preface of the MOF Specification.

As the first adopted technologies specified using a metamodeling approach, the UML and MOF establish a rigorous foundation for OMG's metamodel architectures. Future metamodel standards should reuse their core semantics and emulate their systematic approach to architecture alignment.

Architectural Alignment of UML, MOF, and CORBA

Introduction

This section explains the architectural alignment of the OA&D Facility (OA&DF) metamodel and the MOF meta-metamodel, and their relationships to the OMA and CORBA object models. When discussing specific models, MOF corresponds to the MOF meta-metamodel also referred to as the MOF Model. The UML is used to refer to the proposed OA&DF metamodel.

As yet, there is not an MOF meta-metamodel standard or an OA&D metamodel standard. However, since each of these specifications has been unified, a proactive approach has been taken towards architectural alignment. Considerable structure sharing between the two specifications has been accomplished. As the OA&DF and MOF technologies evolve, additional alignment work will be addressed by standard OMG processes such as those for Revision Task Forces and subsequent RFPs.

The MOF and OA&DF alignment work has focused on aligning the metamodels and applying the MOF IDL Mapping for generating the CORBA IDL for both the MOF and UML models. This was accomplished by defining the MOF and UML models using the MOF and by generating the IDL interfaces based on the MOF specification. Note that both the MOF and OADF specifications use the UML notation for graphically defining the models.

In terms of abstraction levels and the kinds of meta-metaobjects used, the UML and MOF meta-metamodels are well aligned. There are significant advantages in aligning the OA&DF meta-metamodel with the MOF meta-metamodel. In the case of the MOF,

meta-metamodel alignment facilitates interoperability between the OA&DF and the MOF. An example of OA&DF-MOF interoperability is the use of an MOF-compliant repository to store an OA&DF object model.

Alignment of the UML, MOF, and CORBA paves the way for future extensibility of CORBA in key areas such as richer semantics, relationships, and constraints. Likewise the longer-term benefits to UML and MOF include better recognition and addressing of distributed computing issues in developing CORBA-compliant systems.

Motivation

The primary reason for aligning the OA&DF metamodel with the MOF meta-metamodel is to facilitate interoperability between the two facilities using CORBA IDL. When considering interoperability between the OA&DF and the MOF, it is important to consider the difference in scope between the facilities. The MOF goal is to allow interoperability across the application development cycle by supporting the definition of multiple meta models, whereas the OA&DF focuses on supporting the definition of a single OA&D metamodel. An example of OA&DF-MOF interoperability is the use of an MOF-compliant repository to store and interchange OA&DF object models.

The key motivation to align the MOF and OA&DF with CORBA is to address the requirement of aligning with CORBA and between the two facilities. In addition, the MOF and OA&DF (especially the UML) specifications signify years of modeling and metamodeling experience that are being integrated. As such, some of the key concepts in the UML and MOF are potential candidates to evolve the OMG Core object model and CORBA IDL in the future.

Approach

The UML and MOF are based on a four-layer metamodel architecture, where the MOF meta-metamodel is the meta-metamodel for the UML metamodel. As a result, the UML metamodel may be considered an instance-of the MOF meta-metamodel. This is sometimes referred to as loose metamodeling, where an M_n level model is an instance of an M_{n+1} level model.

Since the MOF and OA&DF have different scopes, and diverge in the area of relationships, we have not been able to apply strict metamodeling. In strict metamodeling, every element of an M_n level model is an instance of exactly one element of M_{n+1} level model. Consequently, there is not a strict isomorphic mapping between all the MOF meta-metamodel elements and the UML meta-metamodel elements. In principle strict metamodeling is difficult (or sometimes impossible to accomplish) as the complexity of new concepts (for example patterns and frameworks) continues to increase. In any case, using a small set of primitive concepts such as those defined in the MOF it is possible to define arbitrarily complex metamodels.

In spite of this, since the two models were designed to be interoperable, the two metamodels are structurally quite similar. The following sections compare the core MOF and UML modeling concepts, and contrast them with the OMA and CORBA/IDL core object models. The issues related to mapping metaclasses which are not isomorphic (e.g., Association classes) are also discussed.

The following tables are comparison tables that summarize the mappings of similar concepts across the meta-metamodeling layers. Where there is no analog for a concept, it will be identified and discussed in “Issues Related to UML-MOF Mapping” on page -xvi.

Metamodel Comparison

Most of the metaobjects for the UML core metamodel and the MOF core meta-metamodel can be mapped to each other in a straightforward manner. Similarly, these metaobjects can also be mapped to the OMA and CORBA core object models in a direct way.

The following tables compare the core modeling concepts and the data types for these models.

UML Metamodel	MOF Meta-metamodel	OMA Core Object Model CORBA Object Model	CORBA IDL
Association (n-ary)	Association (binary)		
AssociationClass	NA		
AssociationEnd	AssociationEnd		
Attribute	Attribute	Attribute	Attribute
BehavioralFeature	BehavioralFeature		
Class	Class	Class	Interface (as Class)
Classifier	Classifier		
Constraint	Constraint		
DataType	DataType	Data type	Data type
Dependency (class)	/dependsOn (association)		
Exception	Exception		Exception
Feature	Feature		
GeneralizableElement	GeneralizableElement		
Generalization (class)	generalizes (association)	Generalization	Generalization
Interface	Class (as Interface)	Interface	Interface
ModelElement	ModelElement		
NA	Reference		
NA	Constant		Constant

UML Metamodel	MOF Meta-metamodel	OMA Core Object Model CORBA Object Model	CORBA IDL
Namespace	Namespace		~ IR Containers
Operation	Operation	Operation	Operation
Package	Package		Module
Parameter	Parameter	Parameter	Parameter
StructuralFeature	StructuralFeature		
Type (stereotype)	Class (as Type)	Type	Interface (as Type)

UML Metamodel	MOF Meta-metamodel	CORBA Object Model and IDL
AggregationKind	AggregationKind	
Boolean	CORBA Boolean	Boolean
Enumeration	CORBA Enum	Enum
Expression	NameType	
Integer	CORBA Short, Long, Unsigned Short, Unsigned Long, Double, Octet, Float	Short, Long, Unsigned Short, Unsigned Long, Double, Octet, Float
List	List, Set	Sequence
Multiplicity	MultiplicityKind (simpler than UML multiplicity)	
Name	NameKind	Name
OperationDirectionKind	DirectionKind	
dependencies (reified as classes)	DependencyKind (enum)	
ScopeKind	ScopeKind	
String	CORBA String, Char	String, Char
Time	CORBA Time Service Data Types	Time Service Data Types
NA	TypeDef	TypeDef
Uninterpreted	CORBA Any	Any
VisibilityKind	VisibilityKind	

The MOF supports the full range of CORBA data types as well as additional data types defined using the MOF primitive types. UML supports a subset of CORBA data types in its semantic model but mapping to a subset of specific CORBA types is clearly possible.

The following sections discuss issues related to areas where the mapping between metamodels is not direct.

Issues Related to UML-MOF Mapping

In general, the mapping from the UML meta-metamodel to the MOF meta-metamodel is straightforward.

A review of the previous comparison tables indicates that in most cases the mapping from the UML meta-metamodel to the MOF meta-metamodel is direct. In fact, for most of the core constructs there is a structural equivalency in the mapping.

The key differences are due to different usage scenarios of MOF and UML. The MOF needs to be simpler, directly implementable, and provide a set of CORBA interfaces for manipulating meta objects. The UML is used as a general-purpose modeling language, with potentially many implementation targets. These differences are commonly observed in repository, meta-CASE, and modeling-tool implementations. The key differences are:

- The MOF only supports binary associations while UML supports higher-order (also referred to as 'N-ary') associations. This trade-off was made because N-ary relationships are rarely used in meta-modeling and the design goal was to keep the MOF interfaces simpler. We have anticipated extending the MOF to support higher order associations in future.
- Associations in the MOF are limited to simple associations and cannot contain features. Association Classes in UML can contain features (such as attributes). The MOF has been defined to be structurally extensible to full-blown association classes in the future by relaxing this constraint. UML Association Classes are modeled as MOF Classes with well-defined multiplicity constraints to ensure shared lifetime of features owned by the association.
- The MOF supports the concept of a Reference which allows direct navigation from one Classifier to another. The UML metamodel uses the Target AssociationEnd's 'name' property as a 'pseudo-attribute' for the same purpose, but navigates to another classifier through an Association. The concept of reference is widely used in object repositories, as well as object and object-relational databases and optimizes navigation across a metamodel.
- Some concepts such as Generalization, Dependency, and Refinement are reified as classes in UML, but implemented as Associations in the MOF. The MOF does not require the richness of UML in these areas.
- The MOF supports the full set of CORBA data types, whereas the UML meta model does not define data types to this depth. A CORBA-specific implementation of UML is clearly possible by either creating the additional data types needed or by providing appropriate mappings.
- The UML has clearly defined the similarities of the key concepts of Class, Interface, and Type. The MOF and UML both use the Class concept as the most general of these in their respective models. The MOF specification is focused only on specification of meta models and generation of CORBA interfaces; therefore, it does not deal with implementation concepts such as 'Methods.' UML clearly needs

to support these concepts because UML can be used to model conceptual, logical, and implementation models. In this sense, the MOF uses the Class concept to define Types, since MOF Classes do not have any methods, just operations. This is consistent with the definition of UML Type as a stereotype of Class with a constraint that Types cannot contain methods. The MOF Class concept is rich enough to define Interfaces, and in fact the MOF specification itself clearly shows that an MOF Class can be mapped to multiple CORBA Interfaces.

The previous table shows that the mapping of metadatatypes between the meta-metamodels is also straightforward. Here also there are more MOF meta-metaobjects than there are UML meta-metaobjects. The MOF supports the full range of CORBA data types as well as additional data types defined using the MOF primitive types. UML supports a subset of CORBA data types in its semantic model but maps to specific CORBA types in its corresponding interface model.

Relationship to Other Models

A secondary emphasis was placed on the architectural alignment with CDIF and RM-ODP, both of which have influenced the metamodel architectures. CDIF offers many useful concepts for specifying robust stream-based interchange formats. Similarly, ODP provides many useful ideas for specifying model viewpoints.

Resolution of Technical Criteria

Relationship to OMG IDL and CORE '95

OMG IDL is used to specify all the interfaces in the Meta Object Facility. The MOF itself is of course manipulated using standard CORBA interfaces.

The OMG core object model describes how objects interact; therefore, it is an interface or interaction model. No specific implementation is implied. The OMG object model is not intended to be a metamodel (as described in the OMA). The CORBA object model is a concrete model with the goal of specifying a mechanism for portability and distributed object interoperability. The MOF does define a meta-metamodel (for simplicity, we refer to this as the MOF model). The purpose of the MOF model is to enable the definition and manipulation of metamodels in various domains, with the initial focus being on object analysis and design metamodels. The MOF can be used to specify the OMG object model, which it can treat as a meta-model. Likewise, because the MOF defines a set of CORBA compliant interfaces, these interfaces conform to the CORBA object model. The MOF can be used to specify additional semantics (relationships, constraints) that are implied (or expressed in text) in the CORBA object model.

The MOF is intended to provide support for richer meta data definition and manipulation in a CORBA environment.

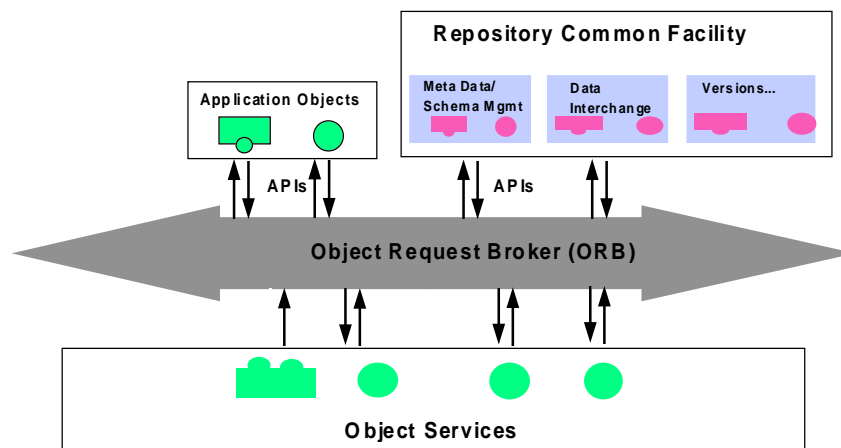
Positioning within the Common Facilities Architecture

The Repository Common Facility is positioned within the Information Management Common Facility and is composed of a number of common facilities and object services, including the MOF and the Change Management Facility.

The standardization of the MOF provides a solid foundation for the OMG architecture in moving toward a unifying architecture for defining and managing meta data in distributed environments. illustrates the positioning of the MOF as a key component of the CORBA architecture as well as within the Common Facilities Architecture.

The Meta Data/Schema Management facility in the figure corresponds to the MOF described in this specification. Note that the MOF interfaces are initially targeted to support the manipulation of OA&D metamodels; however, broader use in areas such as data warehouse management and business object management is expected.

OMG Architecture and the Repository Common Facility



© 1995, Object Management Group, Inc.

OMG Architecture and the Repository Common Facility

Federation of Object Schemas

The MOF allows the definition of metamodels that are potentially domain independent and architecture neutral. The metamodels registered in the MOF can correspond to schemas that are in fact federated. The MOF uses the federation capabilities provided by services such as the Trader Service and the Object Transaction Service. The schema information itself is accessed using the interfaces specified in the MOF. The services of the MOF can be used to define integrated metamodels that correspond to federated schemas. This area is still emerging, and interfaces specifically designed to address federation beyond the services already provided by CORBA (Object Transaction Service and Object Trader service) are not specified in this specification.

Conformance to the MOF

The MOF specification has two conformance points: 1) MOF Model and Interfaces and 2) OMG IDL Generation.

MOF Model and Interfaces

The MOF Model and IDL is the first compliance point. This has the following components:

- The MOF Model and the interfaces of the “Model” module which are defined in Chapter 3. (Additionally, Section B.1, “MOF Model IDL,” on page B-1 gives the consolidated IDL for the “Model” module.)
- The semantics of the “Model” module which are defined by elaborating the MOF to IDL mapping’s semantic specifications in Chapter 5 for the MOF Model.
- The interfaces and semantics of the “Reflective” module which are defined in Chapter 6. (Additionally, Section B.2, “Reflective IDL,” on page B-22 gives the consolidated IDL for the “Reflective” module.)

Document Summary

In addition to this Preface, the *MOF Specification* contains the following chapters:

Chapter 1, Overview, provides several scenarios of domains where MOF is expected to be used. The scenarios discussed include software development, type management, information management, and data warehouse management.

Chapter 2, Model and Interfaces, is the main chapter of the specification. Each of the MOF classes, associations, and data types are fully described along with the IDL interfaces. The semantics of the MOF are defined using a combination of UML notation, textual description, and constraints using the Object Constraint Language (OCL). The MOF interfaces are used to manipulate meta models and meta model constructs in a CORBA environment.

Chapter 3, MOF Packages, defines general purpose interfaces used by all MOF objects to enable self discovery and general purpose manipulation of MOF objects. These interfaces can be used in addition to, or instead of, specific interfaces defined in the MOF Model chapter. These interfaces apply to MOF objects defined in this specification as well as MOF objects defined using the MOF in additional OMG specifications, such as the CORBA interfaces in the UML specification.

Chapter 4, MOF Semantic Details, provides a more comprehensive explanation of the fundamental modeling concepts in the MOF, detailed semantics including structural and behavioral constraints applied to more complex MOF concepts, such as Associations and Packages. This chapter also describes MOF Extensibility Mechanisms and the use of MOF in distributed repository environments.

Chapter 5, MOF to IDL Mapping, defines the generation of CORBA IDL from the descriptions held in the MOF. IDL generation capability is intended to ensure that various metamodels defined using the MOF have consistent IDL. Given an MOF compliant meta-model (such as UML), this portion of the specification can be used to generate a concrete IDL for UML.

Note – This specification does not require that the MOF can derive the meta-object IDL automatically.

Appendix A, MOF IDL Summary, summarizes the complete set of IDL specifications for the MOF and is provided in a format that can be easily processed by IDL compilers.

Appendix B, MODL Description of the MOF, summarizes the Meta Object Definition Language (MODL) textual description of the MOF which was used to generate the CORBA IDL for the MOF.

Appendix C, MOF Implementation Requirements, suggests approaches and principles that can be used to support MOF interoperability between implementation of MOF from multiple vendors. Areas addressed include Model interoperability as well as client tool interoperability when using multiple MOF implementations.

Appendix D, Implementation Requirements, includes vendor boundaries as well as limited implementation requirements.

Appendix E, Future Directions of the MOF, summarizes potential areas of future work related to the MOF based on feedback of MOF submitters and reviewers.

Acknowledgments

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

- BEA Systems, Inc.
- Cooperative Research Centre for Distributed Systems Technology (DSTC)
- Data Access
- Digital Equipment Corporation
- Electronic Data Systems
- Hewlett-Packard Company
- IBM Corporation
- International Computers Limited
- MicroFocus
- Objectivity Inc.
- Oracle Corporation
- Platinum Technology Inc.
- Rational Software Corporation
- System Software Associates
- Unisys Corporation

Contents

This chapter contains the following topics.

Topic	Page
“Overview”	1-1
“Software Development Scenarios”	1-2
“Type Management Scenarios”	1-4
“Information Management Scenarios”	1-6
“Data Warehouse Management Scenarios”	1-6

1.1 Overview

The MOF is intended to support a wide range of usage patterns and applications. To understand the possible usage patterns for the MOF, the first thing one needs to understand is the two distinct viewpoints for the MOF:

1. Modeling viewpoint: The designer's viewpoint, looking "down" the meta levels. From the modeling viewpoint, the MOF is used to define an information model for a particular domain of interest. This definition is then used to drive subsequent software design and/or implementation steps for software connected with the information model.
2. Data viewpoint: The programmer's viewpoint, looking at the current meta-level, and possibly looking up at the higher meta-levels. From the data viewpoint, the MOF (or more accurately, a product of the MOF) is used to apply the OMA-based

distributed computing paradigm to manage information corresponding to a given information model. In this mode, it is possible for a CORBA client to obtain the information model descriptions and to use them to support reflection.

The second thing one needs to realize is that this MOF specification is intended to provide an open-ended information modeling capability. The specification defines a core MOF model that includes a relatively small, though not minimal, set of constructs for object-oriented information modeling. The MOF model can be extended by inheritance and composition to define a richer information model that supports additional constructs. Alternatively, the MOF model can be used as a model for defining information models. This feature allows the designer to define information models that differ from the philosophy or details of the MOF model. In this context, the MOF Model is referred to as a meta-metamodel because it is being used to define metamodels such as the UML.

Finally, one needs to understand the purpose and the limitations of the MOF model to the CORBA IDL mapping defined by this specification. The prime purpose of the mapping is to define CORBA interfaces for information models defined in terms of the MOF model¹ using standard interfaces and interoperable semantics. These interfaces allow a client to create, access, and update information described by the model, with the expectation that the information will be managed in a way that maintains the structural and logical consistency constraints specified in the information model definition.

While we anticipate that some vendors will supply tools (for example, IDL generators, server generators, and so on) to support the development of software conforming to the mapping, provision of these tools is not a requirement of this specification. The second limitation is that the mapping is only intended to support the MOF model itself; that is, it does not support extensions to the metamodel or to other unconnected information models. Furthermore, since the IDL mapping is not itself modeled in the MOF, there can be no standardized support for extending the mapping or defining new mappings. Finally, the IDL mapping in this specification supports only CORBA IDL. Mappings from the MOF model to other interface definition languages are certainly feasible, as are direct mappings to programming languages or data definition languages. However, these mappings are beyond the scope of the first version of the MOF specification.

1.2 Software Development Scenarios

Initially, one of the most likely applications of the MOF will be to support the development of distributed object-oriented software from high-level models. Such a software development system would typically consist of a repository service for storing the computer representations of models and a collection of associated tools. The latter would allow the programmers and designers to input the models, and would assist in the process of translating these models into software implementations.

1. Both extensions to the MOF meta-model that are expressible in the meta-model itself, and unconnected information models expressed using the MOF meta-model.

In the simple case, the repository service could be an implementation of the MOF model interfaces. This service would be accompanied by tools (for example, compilers or graphical editors) that allow the designer to input information models using a human readable notation for the MOF model. Assuming that the target for software development is CORBA based, the system would include an IDL generator that implements the standard MOF model-to-CORBA IDL mapping.

The usage scenario for this repository service would be along the following lines:

1. The programmer uses the input tools provided by the system to define an object-oriented information model using the notation provided.
2. When the design is complete, the programmer runs the IDL generator to translate the model into CORBA IDL.
3. The programmer examines the IDL, repeating steps 1 and 2 to refine the model as required.
4. The programmer then implements the generated IDL to produce a target object server, and implement the applications that use the object server.

The functionality of the development suite described above can be expanded in a variety of ways. We can:

- Add generator tools to automatically produce the skeleton of an object server corresponding to the generated IDL. Depending on the sophistication of the tool, this skeleton might include code for the query and update operations prescribed by the IDL mapping, and code to check the constraints on the information model.
- Add generator tools to produce automatically stereotypical applications such as scripting tools and GUI-based browsers.
- Extend the repository service to store the specifications and/or implementation code for target server and application functionality that cannot be expressed in the MOF model.

While the MOF model is a powerful modeling language for expressing a range of information models, it is not intended to be the ultimate modeling language. Instead, one intended use of the MOF is as a tool for designing and implementing more sophisticated modeling systems. The following example illustrates how the MOF might be used to construct a software development system centered around a hypothetical "Universal Design Language" (UDL).

Many parallels can be drawn between the hypothetical UDL discussed below and the draft OA&DF UML proposal in that UML is designed to be a general purpose modeling language for visualizing, designing, and developing component software. The UDL can be thought of as an extension, as well as a refinement, of many of the concepts in the UML. The extensions are mainly in the area of providing sufficient detail to complete the implementation framework technologies and defining additional meta models that address various technology domains such as database management, transaction processing, etc.

The developer of a software development system based on UDL might start by using an MOF Model notation to define a meta-model for UDL. Conceivably, the UDL metamodel could reuse part or all of the MOF Model, though this is not necessarily a good idea². The developer could then use a simple MOF-based development system (along the lines described above) to translate the UDL metamodel into CORBA IDL for a UDL repository, and to provide hand-written or generated software that implements the UDL repository and suitable UDL model input tools.

The hypothetical UDL development system cannot be considered complete without some level of support for the process of creating working code that implements systems described by the UDL models. Depending on the nature of the UDL, this process might involve a number of steps in which the conceptual design is transformed into more concrete designs and, finally, into program source code. A UDL development system might provide a range of tools to assist the target system designer or programmer. These tools would need to be supported by repository functions to store extra design and implementation information, along with information such as version histories, project schedules, and so on, that form the basis of a mature software development process.

In practice, a software development system implemented along these lines would have difficulty meeting the needs of the marketplace. A typical software engineering "shop" will have requirements on both the technical and the process aspects of software engineering that cannot be met by a "one-size-fits-all" development system. The current trend in software development systems is for Universal Repository systems; that is, for highly flexible systems that can be tailored and extended on the fly.

An MOF-based universal repository system would be based around the core of the MOF Model, and a suite of tools for developing target metamodels (for example, the UDL) and their supporting tools. Many of the tools in the universal repository could be reflective; that is, the tools could make use of information from higher meta-levels to allow them to operate across a range of model types. Functionality, such as persistence, replication, version control, and access control would need to be supported uniformly across the entire repository framework.

1.3 Type Management Scenarios

A second area where early use of the MOF is likely is in the representation and management of the various kinds of type information used by the expanding array of CORBA infrastructure services.

The CORBA Interface Repository (IR) is the most central type-related service in CORBA. The IR serves as a central repository for interface type definitions in a CORBA-based system. The current IR essentially provides access to interface definitions that conform to the implied information model of CORBA IDL. While the

-
2. The MOF meta-model has specific requirements (e.g., model simplicity and support for automatic IDL generation) that are not generally applicable. As a consequence, it is unreasonable to expect the MOF meta-model design to be suitable for all kinds of object modeling.

IR interfaces are tuned fairly well to read-only access, there is no standard update interface and no way to augment the interface definitions in the IR with other relevant information, such as behavioral semantics.

Given a simple MOF-based development environment (as described above), it would be easy to describe the implied information model for CORBA IDL using a notation for the MOF Model. The resulting CORBA IDL model could then be translated into the IDL for an MOF-based replacement for the CORBA IR. While this replacement IR would not be upwards compatible with the existing IR, the fact that it was MOF-based would provide a number of advantages. The MOF-based IR would:

- Support update interfaces.
- Be extensible in the sense that it would be feasible to extend the CORBA IDL model specification by (MOF Model) composition and inheritance. This ability would help smooth the path for future extensions to the CORBA object model.
- Make it easier to federate multiple IR instances and to represent associations between CORBA interface types and other kinds of type information.
- Automatically include links to its own meta-information definition expressed using MOF meta-objects.

Other candidates for use of MOF-based technology among existing and forthcoming infrastructure services include:

- **Trader:** The CORBA trader service maintains a database of "service offers" from services in a CORBA-based distributed environment. These offers have associated service types that are represented using the `CosTradingRepos::ServiceTypeRepository` interface. (A trader service type is a tuple consisting of a type name, an interface type, and a set of named property types. Service types can be defined as subtypes of other service types.)
- **Notification:** At least one initial submission for the forthcoming Notification service includes the notion of an event type. (An event type is a tuple consisting of a type name, a set of named property types, and a set of supertypes.)

In both cases, an MOF-based type repository would have the advantages listed previously for the MOF-based Implementation Repository.

Looking to the future, there are a number of other possible uses for MOF-based type repositories in infrastructure services. For example:

- **Service interface bridges:** As CORBA matures and there is large-scale deployment as part of enterprise-wide computing infrastructures, it will become increasingly necessary to cope with legacy CORBA objects; that is, with objects that provide or use out-of-date service interfaces. In situations where statically deployed object wrappers are not a good solution, one alternative is to provide an ORB-level service that can insert an interface bridge between incompatible interfaces at bind time. Such a service would depend on types that describe the available bridges and the mechanisms used to instantiate them.
- **Complex bindings:** RM-ODP supports the idea that bindings between objects in a distributed environment can be far more complex than simple RPC, stream or multicast protocols. RM-ODP defines the notion of a multi-party binding involving

an arbitrary number of objects of various types, in which different objects fill different roles in the binding. A CORBA service to manage complex bindings would be based on formally described binding types that specify the numbers and types of objects filling each role and the allowed interaction patterns (behaviors) for a given binding.

1.4 Information Management Scenarios

The previous sections focused on the use of the MOF to support the software development life-cycle and the type management requirements of CORBA infrastructure services. This section broadens the scope to the more general domain of information management; that is, the design, implementation, and management of large bodies of more or less structured information.

First, note that some of the ideas outlined above carry over to the information management domain. In some cases, it may be appropriate to define the information model (that is, the database schema) for the application of interest directly using the MOF Model. In this case, the technology described previously can be used to automate the production of CORBA-based servers to store the information and applications to use it. In other situations, the MOF Model can be used to define a metamodel suitable for defining information models for the domain of interest; for example, a metamodel for describing relational database schemas. Then a development environment can be designed and implemented using MOF-based technology that supports the generation of CORBA-based data servers and applications from information models.

In addition, the MOF potentially offers significant benefits for large-scale information systems by allowing such a system to make meta-information available at run-time. Some illustrative examples follow.

Information discovery: The World-Wide Web contains a vast amount of useful (and useless) information on any topic imaginable. However, this information is largely inaccessible. In the absence of other solutions, current generation web indexing systems or search engines must rely on simple word matching. Unless the user frames queries carefully, the number of "hits" returned by a search engine are overwhelming. Furthermore, it is now apparent that even the largest search engines cannot keep pace with the Web's rate of growth.

In the absence of software that can "understand" English text, the approach most likely to succeed is to build databases of meta-data that describe web pages. If this meta-data is represented using MOF-based technology and an agreed base metamodel for the meta-data, the framework can support local meta-data extensions through judicious use of MOF-supported reflection. In addition, because the meta-data framework is defined in the MOF context, it can be accessible to a larger class of generic tools.

1.5 Data Warehouse Management Scenarios

Data warehousing is a recent development in enterprise-scale information management. The data warehouse technique recognizes that it is impractical to manage the information of an enterprise as a unified logical database. Instead, this technique

extracts information from logically- and physically-distinct databases, integrates the information, and stores it in a large-scale "warehouse" database that allows read-only access to possibly non-current data. The extraction and integration processes depend on a database administrator creating a mapping from the schemas for the individual databases to the schema of the warehouse. If the meta-information for the various databases is represented using MOF-based technology, then it should be possible to create sophisticated tools to assist the database administrator in this process.

Meta data is often described as the "heart and soul" of the data warehouse environment. The MOF can be used to automate meta data management of data warehouses. Current meta data repositories that manage data warehouses often use static meta data using batch file-based meta data exchange mechanisms. We expect the use of MOF- and standard CORBA-based event and messaging mechanisms and mobile agent technology (also being standardized by OMG) to drive a new generation of data warehouse management tools and systems that are more dynamic. These tools will enable customers to react in a timelier manner to changing data access patterns and newly discovered patterns, which is the focus of data mining and information discovery systems.

The MOF interfaces and the MOF Model can be used to define specific metamodels for database, data warehouse, model transformation, and warehouse management domains. The integration between these models in a run time data warehouse and the development environment (which has data models) and UML based object models (which describes the corporate data models and operational databases) is a typical use of an MOF. The traceability across these environments is enabled by defining an impact analysis metamodel which builds on the rich model of relationships supported by the MOF.

Contents

This chapter contains the following topics.

Topic	Page
“Overview”	2-1
“Meta-data Architectures”	2-1
“The MOF Model - Meta-modeling Constructs”	2-5
“Meta-models and Mappings”	2-21

2.1 Overview

The Meta Object Facility is a large specification. This chapter aims to make the MOF specification easier to read by providing a conceptual overview of the MOF.

The chapter starts by explaining the MOF’s conceptual architecture for describing and defining meta-data. The next section introduces the meta-modelling constructs that are used to describe meta-data. This is followed by a section that describes how meta-models are mapped to implementation technologies, including the IDL mapping and XMI.

2.2 Meta-data Architectures

The central theme of the MOF approach to meta-data management is openness. The aim is to provide a framework that supports any kind of meta-data, and that allows new kinds to be added as required. In order to achieve this, the MOF uses a layered meta-data architecture that is based on the traditional four layer meta-modeling architecture

that is popular within standards communities such as ISO and CDIF. The key feature of this architecture is a meta-meta-modeling layer that provides a common language that ties together the meta-models and models.

The MOF meta-data architecture is typically (though not exclusively) used as a four layer framework. The MOF Model, which corresponds to the meta-meta-model in a traditional four layer meta-modeling architecture, is an object modeling language that is closely related to UML. The MOF Model is used to define the structure and semantics of general or domain specific meta-models (i.e., schemas for the meta-data of interest). While the MOF Model is object-oriented, it is equally well suited to defining object oriented meta-models, more traditional (e.g. Relational, Entity-Relationship) meta-models, and even simpler ones.

The traditional four layer meta-data architecture is briefly described below. This is followed by a description of how this maps onto the MOF meta-data architecture.

2.2.1 Four Layer Meta-data Architectures

The traditional framework for meta-modeling is based on an architecture with four layers. These layers are conventionally described as follows:

- The user object layer is comprised of the information that we wish to describe. This information is typically referred to as “data.”
- The model layer is comprised of the meta-data that describes information. Meta-data is informally aggregated as models.
- The meta-model layer is comprised of the descriptions (i.e., meta-meta-data) that define the structure and semantics of meta-data. Meta-meta-data is informally aggregated as meta-models. A meta-model can also be thought of as a “language” for describing different kinds of data.
- The meta-meta-model layer is comprised of the description of the structure and semantics of meta-meta-data. In other words, it is the “language” for defining different kinds of meta-data.

The traditional framework is illustrated in Figure 2-1 on page 2-3. This particular example shows how the meta-data for simple records (i.e., “StockQuote” instances) might be represented. The layers are populated as follows:

- The information layer includes some illustrative StockQuote instances.
- The model level includes the meta-data that represents the record type for “StockQuote” instances. The record type has a name (“StockQuote”) and two fields, each of which also has a name and a type. This type will typically be part of some larger scale data schema (not shown here).
- The meta-model level defines what it means to be a record type. The meta-Class for Record is shown as having two meta-Attributes, the first defining the Record’s name, and the second defining its fields. The Meta-Class for a Field (not shown in full) would similarly define the meta-Attributes for the field name and type.

- The meta-meta-model level is typically hard-wired, and defines the machinery that supports the meta-data framework’s meta-modelling constructs; e.g. meta-Classes and meta-Attributes.

Note – In theory, it is possible to add a meta-meta-meta-model and so on. However, for the purposes of the MOF, this 4-layer model suffices.

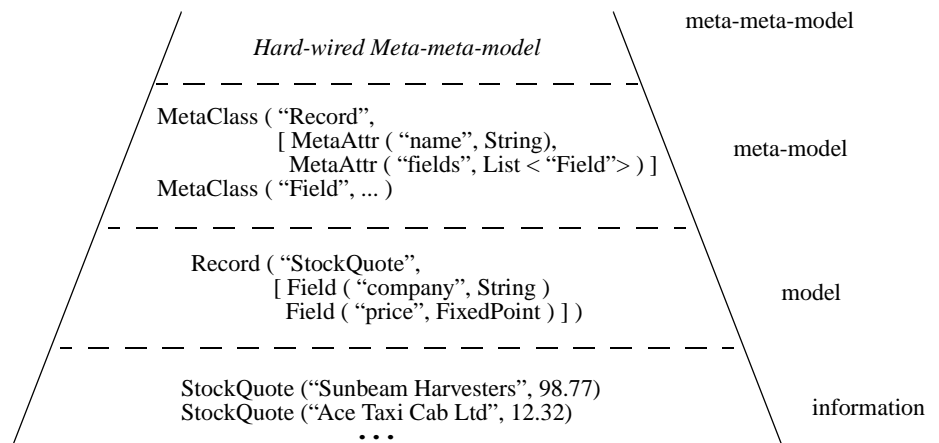


Figure 2-1 Four Layer Meta-data Architecture

While the diagram above shows only one model and one meta-model, the primary aim of having four meta- layers is to support multiple models and meta-models. Just as the model that defines the “StockQuote” type describes many StockQuote instances at the information level, the meta-model that defines “Record” and “Field” can describe many record types at the model level. Similarly, the meta-meta-model level can describe many other meta-models that in turn represent other kinds of meta-data.

The four layer meta-data architecture has a number of advantages:

- Assuming that the meta-meta-model is rich enough, it can support most if not all kinds of meta-information imaginable.
- It potentially allows different kinds of meta-data to be related. (This depends on the design of the framework’s meta-meta-model.)
- It potentially allows interchange of both meta-data (models) and meta-meta-data (meta-models). (This presupposes that the parties to the exchange are using the same meta-meta-model.)

2.2.2 The MOF Meta-data Architecture

The MOF meta-data architecture, shown in Figure 2-2, is based on the traditional four layer meta-data architecture described above.

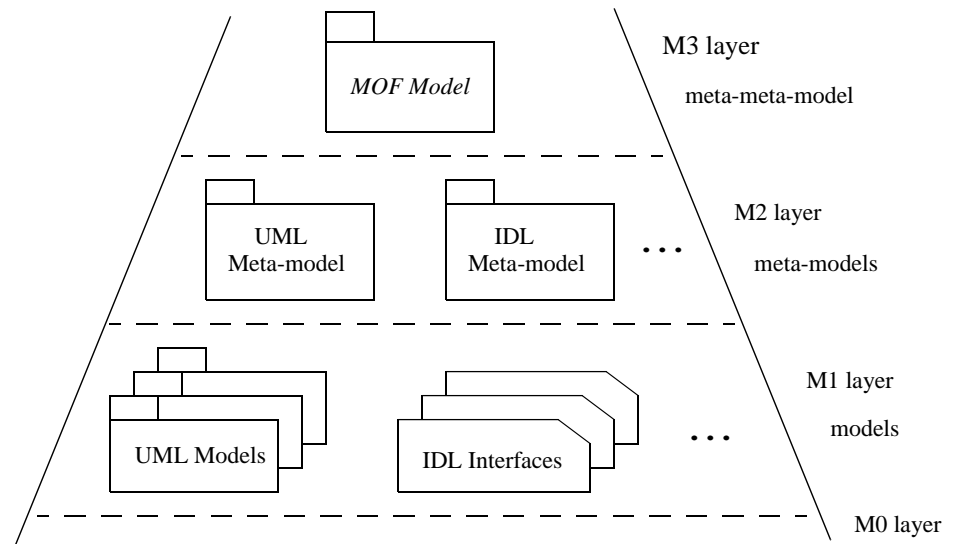


Figure 2-2 MOF Meta-data Architecture

The above diagram shows the MOF meta-data architecture instantiated with meta-models and models for UML and OMG IDL.

The MOF meta-data architecture has a few important features that distinguish it from earlier meta-modeling architectures:

- The MOF Model is object-oriented, supporting meta-modeling constructs that are aligned with (though a bit simpler than) UML's object modeling constructs. Hence, the diagram above uses UML style Package icons to denote MOF-based meta-models as well as UML models.
- The MOF Model is self-describing. In other words, the MOF Model is formally defined using its own meta-modeling constructs. Hence, the MOF Model is also denoted by a UML style Package icon.

The self-defining nature of the MOF Model has some important consequences:

- It helps to validate the MOF's meta-modeling constructs. Since the MOF Model can describe itself, it should be adequate for describing other meta-models of similar complexity.
- It allows the MOF's interfaces and behaviour to be defined by applying the MOF IDL mapping to the MOF Model. This provides uniformity of semantics between computational objects that represent models and meta-models. It also means that when a new technology mapping is defined, the APIs for managing meta-models in that context are implicitly defined as well.
- It provides an architectural basis for future extensions and modifications to the MOF Model.

2.2.3 MOF Meta-modeling Terminology

There is enormous scope for confusion if standard meta-modeling terminology is used in the MOF specification. To avoid this and to make it easier to read, we have opted to simplify the terminology. Some particular points of confusion are as follows:

- The number of MOF meta-levels is not fixed. Since meta-levels are conventionally named upwards from the “information” layer, the meta-level of the top of MOF meta-data framework can vary.
- There are a number of object modeling concepts that appear at two, three, or even four levels in a well populated MOF meta-data framework. For example, a class in a UML is described by an instance of the class “Class” in the UML meta-model. This is in turn described by an instance of the class “Class” in the MOF Model. Finally, the class “Class” in the MOF Model is described by itself.
- While the “meta-” prefix has a clear meaning in the context of the MOF, evidence suggests that people who encounter it for the first time find it very confusing. This is particularly the case for forms like “meta-meta-” and “meta-meta-meta-”.

To avoid some of this confusion, we generally try to avoid using the “meta-” prefix. In particular, while the core of the MOF is a meta-meta-model (assuming that there are 4 meta- layers), it is referred to as “the MOF Model.” Similarly, rather than using terms like Class, MetaClass, and MetaMetaClass, we use phraseology like “an M1-level instance of an M2-level Class.” Terms like M1-level and M2-level are *relative* labels of the meta- levels. (We assume that the reader can mentally adjust the “meta-ness” to fit the context.)

There are three cases where it is convenient to use the “meta-” prefix as part of MOF terminology:

1. The term “meta-data” is used to refer to data whose purpose is to describe other data.
2. The term “meta-model” is used to refer to a model of some kind of meta-data.
3. The term “meta-object” is used to refer to an abstract or technology specific object that represents meta-data.

In each case, the term is used across all meta-levels and has a deliberately imprecise meaning.

The core modeling concepts in the MOF use terms that are common with UML. For example, an MOF Class corresponds to a UML Class, an MOF Attribute corresponds to a UML Attribute, and an MOF Association corresponds to a UML Association. Note however that the correspondence is not always a direct match. For example, UML Associations may have many AssociationEnds, but MOF Associations must have precisely two.

2.3 The MOF Model - Meta-modeling Constructs

This section introduces the MOF’s core meta-modeling constructs (i.e., the MOF’s “abstract language”) for defining meta-models.

MOF meta-modeling is primarily about defining information models for meta-data. The MOF uses an object modeling framework that is essentially a subset of the UML core. In a nutshell, the 4 main modeling concepts are:

1. Classes, which model MOF meta-objects.
2. Associations, which model binary relationships between meta-objects.
3. DataTypes, which model other data (e.g. primitive types, external types, etc.).
4. Packages, which modularize the models.

2.3.1 Classes

Classes are type descriptions of “first class instance” MOF meta-objects. Classes defined at the M2 level logically have instances at the M1 level. These instances have object identity, state, and behavior. The state and behavior of the M1 level instances are defined by the M2 level Class in the context of the common information and computational models defined by the MOF specification.

Instances of classes belong to class extents that impact on certain aspects of their behavior. It is possible to enumerate all instances of a class in a class extent (see Section 4.6.2, “Class Extents,” on page 4-10).

Classes can have three kinds of structural features. Attributes and Operations described below and References described in Section 2.3.4, “References,” on page 2-12. Classes can also contain Exceptions, Constants, DataTypes, Constraints, and other elements.

2.3.1.1 Attributes

An Attribute defines a notional slot or value holder, typically in each instance of its Class. An Attribute has the following properties.

Property	Description
name	Unique in the scope of the Attribute’s Class.
type	May be a Class or a DataType.
“isChangeable” flag	Determines whether the client is provided with an explicit operation to set the attribute’s value.
“isDerived” flag	Determines whether the contents of the notional value holder is part of the “explicit state” of a Class instance, or is derived from other state.
“multiplicity” specification	(see “Attribute and Parameter Multiplicities” on page 2-7)

The aggregation properties of an Attribute depend on the Attribute’s type; see Section 2.3.3, “Aggregation,” on page 2-11.

2.3.1.2 Operations

Operations are “hooks” for accessing behavior associated with a Class. Operations do not actually specify the behavior or the methods that implement that behavior. Instead they simply specify the names and type signatures by which the behavior is invoked. Operations have the following properties.

Property	Description
name	Unique in the scope of the Class.
list of positional parameters having the following properties:	
Parameter name:	
Parameter type	may be denoted by a Class or a DataType
Parameter direction of “in,” “out,” or “in out”	determines whether actual arguments are passed from client to server, server to client, or both.
Parameter “multiplicity” specification	see “Attribute and Parameter Multiplicities” on page 2-7
An optional return type.	
A list of Exceptions that can be raised by an invocation.	

2.3.1.3 Attribute and Operation Scoping

Attributes and Operations can be defined as “classifier level” or “instance level.” An instance-level Attribute has a separate value holder for each instance of a Class. By contrast, a classifier-level Attribute has a value holder that is shared by all instances of the Class in its class extent.

Similarly, an instance-level Operation can only be invoked on an instance of a Class and will typically apply to the state of that instance. By contrast, a classifier-level Operation can be invoked independently of any instance, and can apply to any or all instances in the class extent.

2.3.1.4 Attribute and Parameter Multiplicities

An Attribute or Parameter may be optional-valued, single-valued, or multi-valued depending on its multiplicity specification. This consists of three parts:

1. The “lower” and “upper” fields place bounds on the number of elements in the Attribute or Parameter value. The lower bound may be zero and the upper may be “unbounded.”

- A single-valued Attribute or Parameter has lower bound 1 and upper bound 1. An optional-valued Attribute or Parameter has lower bound 0 and upper bound 1. All other cases are called multi-valued parameters (since their upper bound is greater than 1).

Note – Multiplicity bounds are typically notated as one or two numbers, with “*” used to denote unbounded. For example, a UML bounds specification of “1” translates to lower and upper bounds of 1, and “2..*” translates to a lower bound of 2 and no upper bound.

2. The “is_ordered” flag says whether the order of values in a holder has semantic significance. For example, if an Attribute is ordered, the order of the individual values in an instance of the Attribute will be preserved.
3. The “is_unique” flag says whether instances with equal value are allowed in the given Attribute or Parameter. The meaning of “equal value” depends on the base type of the Attribute or Parameter. See Section 4.2.1, “Semantics of Equality for MOF Values,” on page 4-3, and Section 5.3.1, “Value Types and Equality in the IDL Mapping,” on page 5-6 for additional information.

Note – The bounds and uniqueness parts of a multiplicity specification can give rise to runtime “structural checks” (see “Structural Consistency” on page 2-19). By contrast, orderedness does not imply any runtime checking.

2.3.1.5 *Class Generalization*

The MOF allows Classes to inherit from one or more other Classes. Following the lead of UML, the MOF Model uses the verb “to generalize” to describe the inheritance relationship (i.e., a super-Class generalizes a sub-Class).

The meaning of MOF Class generalization is similar to generalization in UML and to interface inheritance in CORBA IDL. The sub-Class inherits all of the contents of its super-Classes (i.e., all of the super-Classes Attributes, Operations and References, and all nested DataTypes, Exceptions and Constants). Any explicit Constraints that apply to a super-Class and any implicit behavior for the super-Class apply equally to the sub-Class. At the M1 level, an instance of an M2-level Class is type substitutable for instances of its M2-level super-Classes.

The MOF places restrictions on generalization to ensure that it is meaningful and that it can be mapped onto a range of implementation technologies:

- A Class cannot generalize itself, either directly or indirectly.
- A Class cannot generalize another Class if the sub-Class contains a model element with the same name as a model element contained or inherited by the super-Class (i.e., no over-riding is allowed).

- When a Class has multiple super-Classes, no model elements contained or inherited by the super-Classes can have the same name. There is an exception (analogous to the “diamond rule” in CORBA IDL) that allows the super-Classes to inherit names from a common ancestor Class.

Note – It is also possible to use Tags to specify that the interfaces generated for a Class inherits from pre-existing interfaces.

2.3.1.6 *Abstract Classes*

A Class may be defined as “abstract.” An abstract Class is used solely for the purpose of inheritance. No meta-objects can ever exist whose most-derived type corresponds to an abstract Class.

Note – The MOF uses “abstract Class” in the same sense as UML, and also Java and many other object oriented programming languages. Specifying an MOF Class as “abstract” does not say how instances are transmitted. In particular, the use of the term “abstract class” has no relationship to the IDL keyword “abstract” introduced by the Objects-by-value specification.

2.3.1.7 *Leaf and Root Classes*

A Class may be defined as a “leaf” or “root” Class. Declaring a Class as a leaf prevents the creation of any sub-Classes. Declaring a Class as a root prevents the declaration of any super-Classes.

2.3.2 *Associations*

Associations are the MOF Model’s primary construct for expressing the relationships in a meta-model. At the M1 level, an M2 level MOF Association defines relationships (links) between pairs of instances of Classes. Conceptually, these links do not have object identity, and therefore cannot have Attributes or Operations.

2.3.2.1 *Association Ends*

Each MOF Association contains precisely two Association Ends describing the two ends of links. The Association Ends define the following properties..

Property	Description
A name for the end	This is unique within the Association.
A type for the end	This must be a Class.
Multiplicity specification	See “Association End Multiplicities” on page 2-10.

Property	Description
An aggregation specification	See “Association Aggregation” on page 2-12.
A “navigability” setting	Controls whether References can be defined for the end (see Section 2.3.4, “References,” on page 2-12).
A “changeability” setting	Determines whether this end of a link can be updated “in place.”

2.3.2.2 Association End Multiplicities

Each Association End has a multiplicity specification. While these are conceptually similar to Attribute and Operation multiplicities, there are some important differences:

- An Association End multiplicity does not apply to the entire link set. Instead, it applies to projections of the link set for the possible values of the “other” end of a link. See Figure 2-3.
- Since duplicate links are disallowed in M1-level link sets, “is_unique” is implicitly TRUE. The check for duplicate links is based on equality of the instances that they connect; see Section 4.2.1, “Semantics of Equality for MOF Values,” on page 4-3.

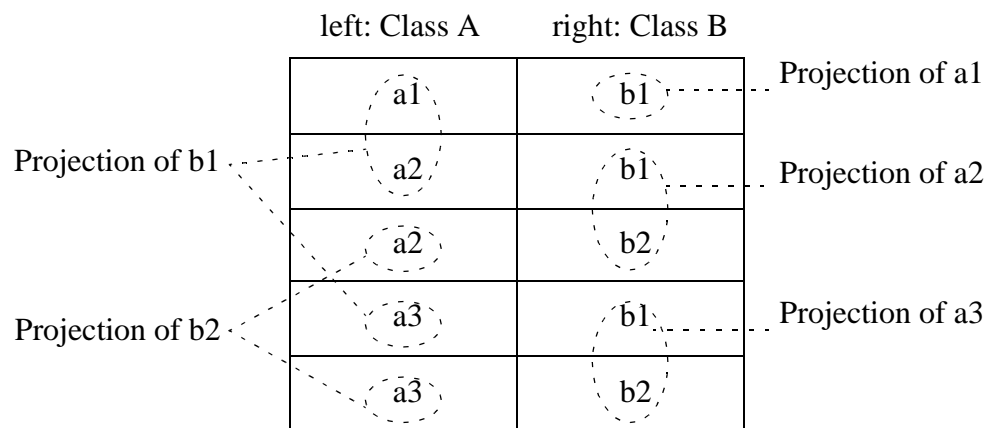


Figure 2-3 The projections of a link set

Figure 2-3 shows a link set for an Association with an AssociationEnd named “left” whose Class is A, and a second named “right” whose Class is B. Instances of A are shown as “a1,” “a2,” and “a3” and “b1” and “b2” are instances of B. In this example with five links, the projection of “a1” is the collection {b1}, and the projection of “b1” is the collection {a1, a2, a3}. If there is another B instance (say “b3”) with no corresponding links, the projection of that b3 is an empty collection.

The “lower” and “upper” bounds of an Association End constrain the number of instances in a projection. For example, if the “left” End of the Association has a bounds “0..3”, then the projection of the link set for any extant instance of B must contain between zero and three instances of A.

The “is_ordered” flag for the Association End determines whether the projections from the other End have an ordering. The MOF Model only allows one of an Association’s two Association Ends to be marked as “ordered.”

In the above example, this could say whether order of the elements of the projection of “b1” is significant (i.e., whether {a1, a2, a3} is a set or a unique list).

2.3.3 Aggregation

In an MOF meta-model Classes and DataTypes can be related to other Classes using Associations or Attributes. In both cases, aspects of the behavior of the relationships can be described as aggregation semantics.

2.3.3.1 Aggregation Semantics

The MOF supports two kinds of aggregation for relationships between instances (i.e., “composite” and “non-aggregate”). A third aggregation semantic - “shared” - is not supported in this version of the MOF specification.

A non-aggregate relationship is a (conceptually) loose binding between instances with the following properties:

- There are no special restrictions on the multiplicity of the relationships.
- There are no special restrictions on the origin of the instances in the relationships.
- The relationships do not impact on the lifecycle semantics of related instances. In particular, deletion of an instance does not cause the deletion of related instances.

By contrast, a composite relationship is a (conceptually) stronger binding between instances with the following properties:

- A composite relationship is asymmetrical, with one end denoting the “composite” or “whole” in the relationship and the other one denoting the “components” or “parts.”
- An instance cannot be a component of more than one composite at a time, under any composite relationship.
- An instance cannot be a component of itself, its components, its components’ components and so on under any composite relationship.
- When a “composite” instance is deleted, all of its components under any composite relationship are also deleted, and all of the components’ components are deleted and so on.
- The Composition Closure Rule: an instance cannot be a component of an instance from a different package extent (see Section 4.9.2, “The Composition Closure Rule,” on page 4-20).

2.3.3.2 Association Aggregation

The aggregation semantics of an Association are specified explicitly using the “aggregation” Attribute of the AssociationEnds. In the case of a “composite” Association, the “aggregation” Attribute of the “composite” AssociationEnd is set to true and the “aggregation” Attribute of the “component” AssociationEnd is set to false. Also, the multiplicity for the “composite” AssociationEnd is required to be “[0..1]” or “[1..1]” in line with the rule that an instance cannot be a component of multiple composites.

2.3.3.3 Attribute Aggregation

The effective aggregation semantics for an Attribute depend on the type of the Attribute. For example:

- An Attribute whose type is expressed as a DataType has “non-aggregate” semantics.
- An Attribute whose type is expressed as a Class has “composite” semantics.

It is possible to use a DataType to encode the type of a Class. Doing this allows the meta-model to define an Attribute whose value or values are instances of a Class without incurring the overhead of “composite” semantics.

2.3.4 References

The MOF Model provides two constructs for modeling relationships between Classes (i.e., Associations and Attributes). While MOF Associations and Attributes are similar from the information modeling standpoint, they have important differences from the standpoints of their computational models and their corresponding mapped interfaces.

Note – Attributes can also model relationships between Classes and DataTypes, but that is not relevant to this point.

Associations offer a “query-oriented” computational model. The user performs operations on an object that notionally encapsulates a collection of links:

- Advantage: The association objects allow the user to perform “global” queries over all relationships, not just those for a given object.
- Disadvantage: The client operations for accessing and updating relationships tend to be more complex.

Attributes offer a “navigation-oriented” computational model. The user typically performs get and set operations on an attribute.

- Advantage: The get and set style of interfaces are simpler, and tend to be more natural for typical meta-data oriented applications which “traverse” a meta-data graph.
- Disadvantage: Performing a “global” query over a relationship expressed as an Attribute is computationally intensive.

The MOF Model provides an additional kind of Class feature called a Reference that provides an alternative “Attribute like” view of Associations. A Reference is specified by giving the following:

- a name for the Reference in its Class,
- an “exposed” Association End in some Association whose type is this Class or a super-Class of this Class, and
- a “referenced” Association End which is the “other” end of the same Association.

Defining a Reference in a Class causes the resulting interface to contain operations with signatures that are identical to those for an “equivalent” Attribute. However, rather than operating on the values in an attribute slot of a Class instance, these operations access and update the Association, or more precisely a projection of the Association. This is illustrated in UML-like notation in Figure 2-4.




Figure 2-4 An example of a Reference

Figure 2-4 shows a Class called `My_Class_1` that is related to `My_Class_2` by the Association `My_Assoc`. `My_Class_1` has an Attribute called “attr” whose type is Integer. In addition, it has a Reference called “ref” that references “end2” of the Association. This provides an API for “ref” that allows a user to access and update a `My_Class_1` instance’s link to a `My_Class_2` instance using get and set operations.

Note – Strictly speaking, the UML notation in the diagram shows “ref” as a derived attribute of `My_Class_1` with type of `My_Class_2`.

The example above shows a Reference that “exposes” an Association End with a multiplicity of “[1..1]”. References can actually expose ends with any valid multiplicity specification. The resulting Reference operations are similar to those for an Attribute with the same multiplicity. However, since MOF Associations do not allow duplicates, Association Ends and therefore References must always have their multiplicity “is_unique” flag set to true.

There are some of important restrictions on References:

- When the “is_navigable” property of an Association End is false, it is not legal to define a Reference that “references” that Association End.
- An M1 instance of a Class that “references” an Association cannot be used to make a link in an instance of the Association in a different extent. This restriction is described in Section 4.9.1, “The Reference Closure Rule,” on page 4-19.

2.3.5 *DataTypes*

Meta-model definitions often need to use attribute and operation parameter values that have “ordinary” types. The MOF provides the meta-modeling concept of a *DataType* to fill this need.

In general terms, *DataTypes* can be used to represent two kinds of type:

1. Meta-models often need to define types whose values do not have object identity; e.g. integers, strings, enumerations and so on.
2. Meta-models sometimes need to reuse “external” types; i.e. types which are defined in some kind of non-MOF interface specification.

Note – The current MOF specification only states how CORBA data types and (external) CORBA interface types are handled. Support for other type systems is scheduled for inclusion in the next revision.

See Section 3.4.7, “*DataType*,” on page 3-33 for more details on how *DataTypes* are used to express types.

2.3.6 *Packages*

The *Package* is the MOF Model construct for grouping elements into a meta-model. *Packages* serve two purposes.

1. At the M2 level, *Packages* provide a way of partitioning and modularizing the meta-model space. *Packages* can contain most kinds of model element (e.g., other *Packages*, *Classes*, *Associations*, *DataTypes*, *Exceptions*, *Constants* and so on).
2. At the M1 level, *Package* instances act the outermost containers for meta-data. Indirectly, they also define the scope boundaries of *Association* link sets and of “classifier level” *Attributes* and *Operations* on *Class* instances (see Section 4.6.4, “*Package Extents*,” on page 4-11).

The MOF Model provides four mechanisms for meta-model composition and reuse (i.e., generalization, nesting, importing, and clustering). These are described in the following subsections.

2.3.6.1 *Package Generalization*

Packages may be generalized by (inherit from) one or more other *Packages* in a way that is analogous to *Class* generalization described in Section 2.3.1.5, “*Class Generalization*,” on page 2-8. When one *Package* inherits from another, the inheriting (sub-) *Package* acquires all of the meta-model elements belonging to the (super-) *Package* it inherits from. *Package* inheritance is subject to rules that prevent name collision between inherited and locally defined meta-model elements.

At the M1 level, a sub-Package instance has the ability to create and manage its own collections of Class instances and Links. This applies to the Classes and Associations that it defines explicitly, and to those that it acquires by inheritance.

The relationship between instances of the super- and sub-Packages is similar to relationship between instances of super- and sub-Classes:

- A sub-Package instance is type substitutable for instances of its super-Packages (i.e., the sub-Package instance “IS_A” super-Package instance).
- A sub-Package instance does not use or depend on an instance of the super-Package (i.e., there is no “IS_PART_OF” relationship).

Packages may be defined as “root” or “leaf” Packages (with analogous meaning to “root” and “leaf” Classes), but “abstract” Packages are not supported.

2.3.6.2 *Package Nesting*

A Package may contain other Packages, which may in turn contain other Packages. Model elements defined in nested Packages may be strongly coupled to other model elements in the same containment. For example, a Class in a nested Package have a Reference that links it via an Association in its context, or its semantics could be covered by a user-defined Constraint that applies to the enclosing Package.

A nested Package is a component of its enclosing Package. Since, in general, the model elements in a nested Package can be inextricably tied to its context, there are some significant restrictions on how nested Packages can be composed. In particular,

- a nested Package may not generalize or be generalized by other Packages.
- a nested Package may not be imported or clustered by other Packages.

Nested Packages are not directly instantiable. No factory objects or operations are defined for nested Package instances. An M1 level instance of a nested Package can only exist in conjunction with an instance of its containing Package. Conceptually, a nested Package instance is a component of an instance of its containing Package.

Note – The main effect of nesting one Package inside another is to partition the concepts and the namespace of the outer Package. Nesting is not a mechanism for reuse. Indeed when a Package is nested, the options for reusing its contents are curtailed.

2.3.6.3 *Package Importing*

In many situations, the semantics of Package nesting and generalization do not provide the best mechanism for meta-model composition. For example, the meta-modeler may wish to reuse some elements of an existing meta-model and not others. The MOF provides an import mechanism to support this.

A Package may be defined as importing one or more other Packages. When one Package imports another, the importing Package is allowed to make use of elements defined in the imported one Package. As a shorthand, we say that the elements of the imported Package are imported.

Here are some examples of how a Package can reuse imported elements. The importing Package can declare:

- Attributes, Operations, or Exceptions using imported Classes or DataTypes,
- Operations that raise imported Exceptions,
- DataTypes and Constants using imported DataTypes or Constants,
- Classes whose supertypes are imported Classes, and
- Associations for which the types of one or both Association Ends is an imported Class.

At the M1 level, an instance of an importing Package has no explicit relationship with any instances of the Packages that it imports. Unlike a subtype Package, an importing Package does not have the capability to create instances of imported Classes. A client must obtain any imported Class instances it needs via a separate instance of the imported Package.

2.3.6.4 *Package Clustering*

Package clustering is a stronger form of Package import that binds the importing and imported Package into a “cluster.” As with ordinary imports, a Package can cluster a number of other Packages, and can be clustered by a number of other Packages.

An instance of a cluster Package behaves as if the clustered Packages were nested within the Package. That is, the lifecycle of a clustered Package instance is bound to the lifecycle of its cluster Package instance. In particular:

- When the user creates an instance of a cluster Package, an instance of each of its clustered Packages is created automatically.
- The instances of the clustered Packages created above all belong to the same cluster Package extent.
- Deleting a cluster Package instance automatically deletes its clustered Packaged instances, and the clustered Package instances cannot be deleted except as part of the deletion of the cluster Package instance.

However, unlike a nested Package, it is possible to create an independent instance of a clustered Package. Also, in some situations clustered Package instances are not strictly nested.

Note – It is possible to cluster or inherit from Packages that cluster other Packages. The impact of this on M1 level instance relationships is discussed in Section 4.6.4, “Package Extents,” on page 4-11.

In summary, the relationship between the M1 level instances in a Package cluster is that each clustered Package instance is a component of the cluster Package instance. Unlike nested Packages, there is no composite relationship between the M2 level Packages.

2.3.6.5 Summary of Package Composition Constructs

The properties of the four Package composition mechanisms defined by the MOF Model are summarized by Table 2-1.

Table 2-1 Package Composition Constructs

Meta-model Construct	Conceptual Relationship	M2 level Relationship Properties	M1 level Relationship Properties
Nesting	P1 contains P2	$P1 \blacklozenge \longrightarrow P2$	$\underline{P1} \blacklozenge \longrightarrow \underline{P2}$
Generalization / Inheritance	P1 generalizes P2	$P2 - - \Rightarrow P1$	$\underline{P2} \longrightarrow \triangle P1$
Importing	P1 imports P2	$P1 - - - \Rightarrow P2$	none
Clustering	P1 clusters P2	$P1 - - \Rightarrow P2$	$\underline{P1} \diamond \longrightarrow \underline{P2}$ or none

The symbology of the table is based on UML (i.e., a filled diamond means composition, a hollow diamond means aggregation, a hollow triangle means inheritance, and a dotted arrow means “depends on”).

Note that P1 and P2 denote different (though related) things in different columns of the table:

- In column 2, they denote conceptual M2 level Packages in a meta-model.
- In column 3, they denote both the conceptual M2 level Packages, and the objects that represent them in a reified meta-model.
- In column 4, they denote M1 level Package instances (when underlined) or their types.

2.3.7 Constraints and Consistency

The MOF Model constructs described so far allow the meta-modeler to define a meta-data information that comprises nodes (Classes) with attached properties (Attributes / DataTypes) and relationships between nodes (Associations). While the above constructs are sufficient to define an “abstract syntax” consistent of meta-data nodes and links, this syntax typically needs to be augmented with additional consistency rules.

This section describes the MOF Model’s support for consistency rules and model validation.

2.3.7.1 Constraints

The MOF Model defines an element called Constraint that can be used to attach consistency rules to other meta-model components. A Constraint comprises:

- a constraint name,
- a “language” that identifies the language used to express the consistency rules,
- an “expression” in the language that specifies a rule,
- an “evaluation policy” that determines when the rule should be enforced, and
- a set of “constrained elements.”

A Constraint expression is an expression in some language that can be “evaluated” in the context of a meta-model to decide if it is valid. The MOF specification does not define or mandate any particular languages for Constraint expressions, or any particular evaluation mechanisms. Indeed, it is legitimate for Constraints to be expressed in informal language (e.g., English) and for validation to be implemented by ad-hoc programming. However, the Constraints that are part of the MOF Model specification itself are expressed in Object Constraint Language (OCL) as described in the UML specification.

The evaluation policy property of a Constraint determines whether the consistency rule should be enforced immediately or at a later time. Figure 2-5 gives a simple example that will be used to illustrate the need for evaluation policies.

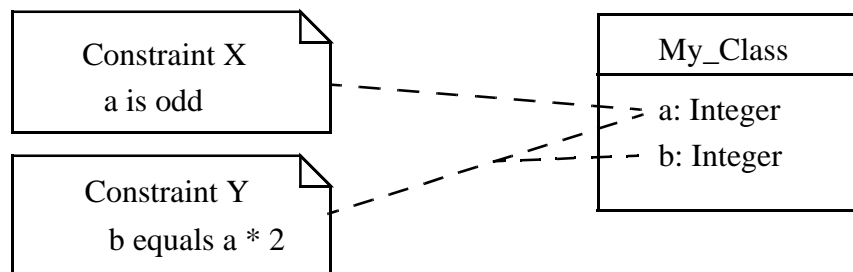


Figure 2-5 Examples of Constraints

In Figure 2-5, Constraint X constrains only Attribute a while Constraint B constrains both Attributes a and b.

It is feasible to check the first Constraint (X: “a is odd” on the Attribute “a”) at any time. It could be checked whenever a value for “a” is supplied (e.g., at instance creation and when “a” is updated). An exception would be raised if the new value for “a” was even. Alternatively, constraint checking could be deferred to a later point (e.g., when the user requests validation of a model).

The second constraint (Y: “b equals a * 2” on both Attributes “a” and “b”) is another matter. If a server enforces Y on every update, the user would never be able to change the values of either “a” or “b”. No matter which order the user invoked the operations, the updates would raise an exception. Instead, enforcement of Y must be deferred until both “a” and “b” have been updated.

Note – The Constraint construct is intended to be used for specifying consistency rules for models rather than for defining the computation behavior of (for example) Operations. It is “bad style” to specify Constraint expressions that have side-effects on the state of a model, not least because it is unspecified when Constraints are evaluated.

2.3.7.2 *Structural Consistency*

As noted previously, an MOF-based meta-model defines an “abstract syntax” for meta-data. Some aspects of the abstract syntax are enforced by the corresponding meta-data server’s IDL. For example, the operation that creates a link for an Association has a type signature that prevents the user from creating a link with the wrong kind of Class instances. However, some aspects of the abstract syntax can only be enforced by runtime structural consistency checks. While most of the structural checks are made immediately, checks for “underflow” often need to be deferred.

It is not practical for a meta-model to specify *a priori* all possible things that can go wrong in a/n MOF-based meta-data server. It is therefore necessary to recognize that a/n MOF server may need to perform a variety of runtime checks that are neither defined or implied by the meta-model. These include additional meta-data validation that is not specified by the meta-model, resource and access control checks, and internal error checking.

2.3.7.3 *Consistency Checking Mechanisms*

The MOF specification provides a lot of latitude for meta-data server implementations in the area of constraint checking or validation.

- Support for checking of Constraints is not mandatory. In particular, there is no requirement to support any particular language for Constraint expressions.
- The set of events (if any) that may trigger deferred checking is not specified. No general APIs are specified for initiating deferred consistency checking.
- Persistence and interchange of meta-data, which is in an inconsistent state may be allowed. (Indeed, this would seem to be a prerequisite for some styles of meta-data acquisition.)
- There are no specified mechanisms for ensuring that validated meta-data remains valid, or that it does not change.

The one aspect of consistency checking that is mandatory is that a meta-data server must implement all structural consistency checks that are labeled as immediate.

2.3.8 Miscellaneous Meta-modeling Constructs

This section describes the remaining significant elements of the MOF Model.

2.3.8.1 Constants

The Constant model element allows the meta-modeler to define simple bindings between a name and a constant value. A Constant simply maps onto a constant declaration in (for example) the IDL produced by the MOF IDL mapping.

2.3.8.2 Exceptions

The Exception model element allows the meta-modeler to declare the signature of an exception that can be raised by an Operation. An Exception simply maps onto (for example) an IDL exception declaration.

2.3.8.3 Tags

The Tag model element is the basis of a mechanism that allows a “pure” MOF meta-model to be extended or modified. A Tag consists of:

- a name that can be used to denote the Tag in its container,
- a “tag id” that denotes the Tag’s kind,
- a collection of zero or more “values” associated with the Tag, and
- the set of other model elements that the Tag is “attached” to.

The meaning of a model element is (notionally) modified by attaching a Tag to it. The Tag’s “tag id” categorizes the intended meaning of the extension or modification. The “values” then further parameterize the meaning.

As a general rule, the definition of values and meanings for “tag id” strings is beyond the scope of the MOF specification. The specification recommends a tag id naming scheme that is designed to minimize the risk of name collision, but use of this scheme is not mandatory; see Section 3.4.23, “Tag,” on page 3-63.

One exception to this is the MOF to IDL Mapping. This defines some standard tag ids that allow a meta-model to influence the IDL mapping; see Section 5.6, “Standard Tags for the IDL Mapping,” on page 5-35 for the complete list. For example:

- “Substitute Name” provides an alternative IDL identifier for an element in a meta-model, and
- “IDL Prefix” allows the meta-modeler to specify the IDL “prefix” for a top-level Package.

2.4 Meta-models and Mappings

The previous sections outlined the overall meta-data architecture for the MOF, and the meta-modeling constructs provided by the MOF Model. This section describes the Mapping approach that is used to instantiate MOF meta-models and meta-data in the context of a given implementation technology.

This section is organized as follows. The first subsection outlines the purpose and structure of MOF Mappings. The next two subsections give high-level overviews of the OMG MOF technology mappings defined to date. The final subsection explains how the standard mappings are applied to the MOF Model to produce the OMG IDL for the MOF Model server and an XML DTD for meta-model interchange.

2.4.1 Abstract and Concrete Mappings

MOF Mappings relate an M2-level meta-model specification to other M2 and M1-level artifacts, as depicted in Figure 2-6.

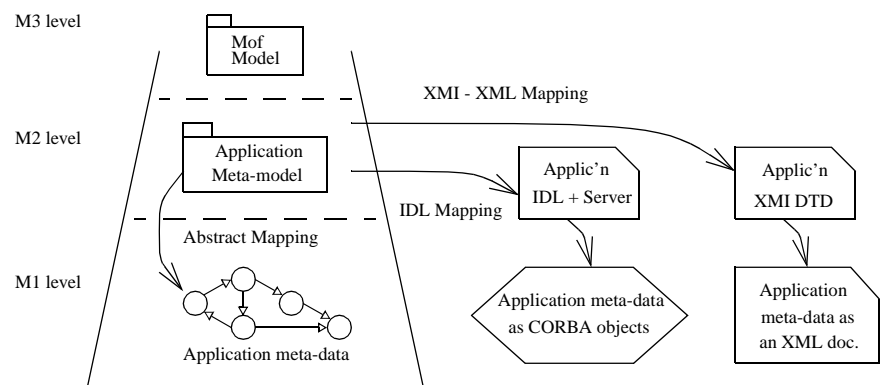


Figure 2-6 The function of MOF Technology Mappings

Figure 2-6 depicts the Mapping derived relationships for an application meta-model as follows:

- The Abstract mapping (defined in “The MOF Abstract Mapping” chapter) fleshes out a MOF meta-model into an abstract information model (i.e., by spelling out the logical structure of the meta-data described by the meta-model).
- The IDL Mapping (Section 2.4.2, “CORBA Meta-data Services - The MOF IDL Mapping,” on page 2-22) produces the standard OMG IDL and associated behavioural semantics for meta-objects that can represent meta-data conforming to the meta-model.
- The XML Mapping (see Section 2.4.3, “Meta-data Interchange - The MOF XML Mapping,” on page 2-22) produces the standard XML DTD for interchanging meta-data conforming to the meta-model.

The Abstract Mapping has two roles: 1) it serves to define the “meaning” of a meta-model, and 2) it provides a point of alignment for current and future MOF technology Mappings.

Since the IDL and XML Mappings are both aligned with the Abstract Mapping there is a mechanical one-to-one correspondence between abstract meta-data and meta-data expressed as XMI documents and as CORBA meta-objects. This correspondence holds for all meta-models. More significantly, it should also hold for any future meta-model Mappings (e.g., to Java or DCOM technology) provided that they are aligned with the Abstract Mapping.

2.4.2 CORBA Meta-data Services - The MOF IDL Mapping

The MOF IDL Mapping produces a specification for a CORBA meta-data service from a MOF meta-model specification. The OMG IDL interfaces and associated behavioral semantics are specified in the “MOF to IDL Mapping” chapter and “The Reflective Module” chapter. These interfaces support creating, updating, and accessing meta-data in the form of CORBA objects, either using “specific” interfaces that are tailored to the meta-model or “generic” interfaces that are meta-model independent.

The MOF IDL Mapping places some additional restrictions in MOF meta-models beyond those set out in the “MOF Model and Interfaces” chapter. See Section 5.5, “Preconditions for IDL Generation,” on page 5-33 for details.

2.4.3 Meta-data Interchange - The MOF XML Mapping

Interchange of MOF-based meta-data is defined in a separate OMG specification. The XMI (XML-based Meta-data Interchange) specification leverages the W3C’s XML (eXtensible Markup Language) technology to support the interchange of meta-data and meta-models between MOF-based and other meta-data repositories.

The XMI specification (ad/98-07-01) has two main parts:

1. The “XML DTD Production Rules” define a uni-directional mapping from a/n MOF meta-model to a XML DTD (Document Type Definition) for meta-data interchange documents.
2. The “XML Document Production Rules” define a bi-directional mapping between an XML document (structured according to the above DTD) and MOF-based meta-data that (implicitly) conforms to the Abstract Mapping.

2.4.4 Mappings of the MOF Model

The XMI specification has been applied to the MOF Model (rendered as a/n MOF meta-model) to produce XML and DTD documents that form an optional compliance point of the MOF specification:

- Appendix A.1, “The MOF Model in XML” contains an XMI rendering of the MOF Model itself, along with the XMI generated DTD for MOF meta-models. This XML document should be viewed as a normative rendering of the MOF Model for interchange purposes.
- Appendix A.2, “The XMI DTD for MOF meta-models” contains the normative XMI DTD for the interchange of MOF meta-models.

Contents

This chapter contains the following topics.

Topic	Page
“Overview”	3-1
“How the MOF Model is Described”	3-2
“The Structure of the MOF Model”	3-11
“MOF Model Classes”	3-15
“MOF Model Associations”	3-65
“MOF Model Data Types”	3-77
“MOF Model Exceptions”	3-83
“MOF Model Constants”	3-85
“MOF Model Constraints”	3-86

3.1 Overview

This chapter describes the model that defines the MOF. The MOF provides a set of modeling elements, including the rules for their use, with which to construct models. Specifically, the MOF modeling elements support development of meta-models. This focus enables the MOF to provide a more domain-specific modeling environment for defining meta-models instead of a general-purpose modeling environment.

A well-designed modeling tool or facility should be based on a meta-model that represents the modeling elements and the rules provided by the tool or facility.

Every meta-model is also a model. If the MOF Model described in this section is the meta-model for the MOF, where is the model for this meta-model? Formally, the MOF is defined in itself; that is, the modeling elements defined in the MOF Model and provided by the MOF are used to define the MOF Model itself. In essence, the MOF Model is its own meta-model. However, this circular definition does not support presentation of the model. Therefore, this specification describes the MOF narratively and through the use of UML notation, tables, and Object Constraint Language (OCL) expressions.

Note that the use of UML notation is a convenience to the designers of the MOF and to the readers of the MOF specification. The semantics of the MOF Model are completely defined in the MOF specification and do not depend on the semantics of any other model. The MOF interfaces used to manipulate meta-models are dependent on CORBA in that these interfaces are specified using CORBA IDL.

A significant amount of the MOF Model syntax and semantics definition is constraint-based. This specification describes the constraint expressions as clearly as possible. In addition, the specification provides a reference to the OCL expression that defines each constraint.

The OCL, which is defined in Object Constraint Language Definition, provides a small set of language elements used to define expressions (see the Unified Modeling Language Specification for additional OCL information). As an expression language, OCL cannot change the state of objects; however, it can express constraints (including invariants, preconditions, and post-conditions). OCL expressions use operations defined in the MOF Model with the attribute `isQuery` set to `TRUE`. (Such operations do not change the state of the object.) To ensure complete specification of constraints, this document provides OCL definitions for MOF-defined operations used in OCL expressions. In addition, to avoid ambiguity or misinterpretation this specification uses OCL to define a few of the most complex concepts of the MOF Model.

The interfaces through which the MOF is utilized are generated from the MOF Model. However, these interfaces do not provide the semantic information necessary to determine the behavior of their operations. Therefore, it is essential to understand the MOF in terms of its model and related semantics, not just its interfaces.

3.2 *How the MOF Model is Described*

This chapter describes the modeling elements that comprise the MOF Model and provide the building blocks for meta-models. Because these elements are formally described with the MOF Model itself, the characteristics used to describe the model are the same characteristics provided by the model.

The following subsections briefly describe the conventions that this specification uses to define the model elements and their characteristics, with a few exceptions noted.

3.2.1 Classes

Classes are the fundamental building blocks of MOF meta-models and the MOF Model. A Class can have three kinds of features; Attributes, References, and Operations. They may inherit from other Classes, and may be related to other Classes by Associations. Classes are presented in detail in Section 4.3, “Semantics of Class Instances,” on page 4-3.

The MOF uses the term Class with a meaning that is similar to that of Class in UML. A/n MOF Class is an abstract specification or classification of meta-objects that includes their state, their interfaces, and (at least informally) behavior. A Class specification is sufficient to allow the generation of concrete interfaces with well defined semantics for managing meta-object state. However, a/n MOF Class specification does not include any methods to implement meta-object behavior.

The Classes that make up the MOF Model are introduced in Section 3.3, “The Structure of the MOF Model,” on page 3-11 and specified in detail in Section 3.4, “MOF Model Classes,” on page 3-15. Each Class is defined in terms of its name(s), its super-Classes, the Classes whose instances it can contain, its attributes, its references, its operations, its constraints, and whether it is abstract or concrete.

Note – Except where stated, the order in which Section 3.4 introduces Classes and their component features is not normative. The normative order is defined in the XMI for the MOF Model which may be found in Appendix A. This order determines the order in which elements appear in the generated IDL, and is in theory significant.

This document uses a hybrid textual and tabular notation to define the important characteristics of each Class in the MOF Model. The notation defines defaults for most characteristics, so that the Class definitions need only explicitly specify characteristics that are different from the default. The following text explains the notation used for Classes and their characteristic.

3.2.1.1 Class Heading

Each Class in the MOF Model is introduced by a second level section heading. The heading defines the standard ModelElement name for the Class. The Classes name on the heading line can be followed by the word “*abstract*” or by a “*substitute_name*” for some mapping. For example, the following:

Chapter 3.4.1 ModelElement *abstract*

introduces a Class called “ModelElement” and defines its Chapter “isAbstract” flag to have the value “true.” On the other hand, the following:

Chapter 3.4.11 Attribute *idl_substitute_name* *“MofAttribute”*

introduces a Class called “Attribute” and defines its substitute name (for the IDL mapping) as “MofAttribute.” The latter information is encoded using a Tag whose “tagId” is “idl_substitute_name” and whose “values” consist of the Any-ized string “MofAttribute.”

Unless stated otherwise each Class in the MOF Model has “isAbstract” set to false, and has no attached Tags.

Note – The MOF uses “abstract Class” in the same sense as UML, and also Java and many other object oriented programming languages. There is no relationship with the IDL keyword “abstract” introduced in CORBA 2.3.

The paragraph or paragraphs following a Class heading give a description of the Class, its purpose, and its meaning.

3.2.1.2 *Superclasses*

The “Superclasses” heading lists the MOF Classes that generalize the Class being described. In the MOF context, generalization is another term for inheritance. Saying that a Class A generalizes a Class B, means the same as saying that Class B inherits from Class A. The sub-Class (B) inherits the contents of the super-Class (A). Multiple inheritance is permitted in the MOF.

This heading is always present, since with the sole exception of ModelElement, all Classes in the MOF Model have super-Classes.

3.2.1.3 *Contained Elements*

Instances of the sub-Classes of Namespace can act as containers of other elements. If present, the “Contained Elements” heading lists the Classes whose instances may be contained by an instance of this container Class. It also gives the index of the MOF Model Constraint that defines the containment rule for the Class. For more details, see Section 3.3.3, “The MOF Model Structure,” on page 3-13. In particular, Table 3-4 on page 3-15 expresses the MOF Class containment rules in a concise form.

If the “Contained Elements” heading is absent, instances of the Class may not contain other instances. This occurs if the Class is an abstract Class (and therefore has no instances), or if the Class is not derived from the Namespace Class.

3.2.1.4 *Attributes*

The “Attributes” heading lists the Attributes for a Class in the MOF Model. Attributes that are inherited from the super-Classes are not listed. If the “Attributes” heading is missing, the Class has no Attributes.

All Attributes defined in the MOF Model have a “visibility” of “public_vis.” All have a “type” that is represented using a DataType, and therefore all have aggregation semantics of “none.” The remaining characteristics of Attributes are defined using the notation described in Table 3-1.

Table 3-1 Notation for Attribute Characteristics

Entry	Description
<i>type:</i>	This entry defines the base type for the Attribute. This is expressed as the name of a DataType defined in Section 3.6, “MOF Model Data Types,” on page 3-93, or as the name of a CORBA data type (e.g., “boolean” or “string”). The base type is represented by the Attribute’s “type.”
<i>multiplicity:</i>	This entry defines the “multiplicity” for the Attribute, consisting of its “lower” and “upper” bounds, an “isOrdered” flag, and an “isUnique” flag. See Section 3.6.5, “MultiplicityType,” on page 3-94, and Section 4.4.2, “Multiplicity,” on page 4-5 for more details. The multiplicity for an Attribute is expressed as follows: <ul style="list-style-type: none"> • The “lower” and “upper” bounds are expressed as “exactly one,” “zero or one,” “zero or more,” and “one or more.” • If the word “ordered” appears, “isOrdered” should be true. If it is absent, “isOrdered” should be false. • If the word “unique” appears, “isUnique” should be true. If it is absent, “isUnique” should be false.
<i>changeable:</i>	This optional entry defines the “isChangeable” flag for the Attribute. If omitted, “isChangeable” is true.
<i>derived from:</i>	This optional entry either describes the derivation of a derived Attribute, or if the entry is present, the Attribute’s “isDerived” flag will be true. If it is absent, the flag will be false.
<i>scope:</i>	This optional entry defines the “scope” of an Attribute as either “instance_level” or “classifier_level.” If the entry is absent, the Attribute’s “scope” is “instance_level.”

3.2.1.5 References

The “References” heading lists the References for a Class in the MOF Model. A Reference connects its containing Class to an Association End belonging to an Association that involves the Class. This allows a client to navigate directly from an instance of the Class to other instance or instances that are related by links in the Association. If the “References” heading is absent, the Class has no References.

A Class involved in an Association may or may not have a corresponding Reference. A Reference means that a client can navigate to instances of the other Class; however, this comes at the cost of some restrictions. In particular, if one or both Classes in an Association have References for that Association, the Reference Closure rule restricts the creation of links between instances in different “extents” (see Section 4.9.1, “The Reference Closure Rule,” on page 4-17).

Note – The modeling of navigation in MOF differs from UML. In UML, mechanisms for navigating links are available when the “isNavigable” flag is true for a given AssociationEnd. In this case, stronger uniqueness constraints on AssociationEnd names mean that they are unique within the namespaces of the Association and all Classes involved and their sub-Classes. This means that the AssociationEnd names uniquely bind to a “navigator” operation in each context in which navigation might be used.

Most characteristics of References in the MOF Model are either common across all References or derived from other information:

- The “visibility” of all References in the MOF Model is “public_vis.”
- The “scope” of all References is “instance_scope.”
- The “type” of all References is the same as the “type” of the AssociationEnd it references.
- The “multiplicity” of all References is the same as the “multiplicity” of the AssociationEnd it references.

The variable characteristics of References are defined or documented using the notation described in Table 3-2.

Table 3-2 Notation for Reference characteristics

Entry	Description
<i>class:</i>	This entry documents the base type of the Reference and is represented as its “type.” Note that the “type” of a Reference must be the same as the “type” of the referenced AssociationEnd.
<i>defined by:</i>	This entry defines the Association and AssociationEnd that the Reference is linked to via a RefersTo link.
<i>multiplicity:</i>	This entry documents the “multiplicity” characteristics for the Reference. These are written the same way as Attribute “multiplicity” characteristics, except that “unique” is omitted because its value is predetermined (see Section 3.2.2.2, “Ends,” on page 3-8). Note the following: <ul style="list-style-type: none"> • the OCL constraints on MultiplicityType and AssociationEnd mean that the “isUnique” field must be “false” if the “upper” bound is 1 and “true” otherwise, and • the “multiplicity” settings for an AssociationEnd and its corresponding Reference(s) must be the same.
<i>changeable:</i>	This optional entry defines the setting of the Reference’s “isChangeable” flag. If the entry is absent, the “isChangeable” flag is true.
<i>inverse:</i>	This optional entry documents the “inverse” Reference for this Reference (i.e., the Reference on the link related Class that allows navigation back to this Reference’s Class). If this entry is absent, the Reference does not have an inverse Reference.

3.2.1.6 Operations

The “Operations” heading lists the Operations for a Class in the MOF Model. If the heading is absent, the Class has no Operations.

All Operations for Classes in the MOF Model have “visibility” of “public_vis.” The remaining characteristics of References are defined using notation described in Table 3-3.

Table 3-3 Notation for Operation Characteristics

Entry	Description
<i>return type:</i>	<p>This optional entry defines the “type” and “multiplicity” of the Operation’s return Parameter (i.e., the one with “direction” of “return_dir”). The “type” is denoted by a name of a Class or DataType in the MOF Model, or a name of a CORBA data type. The “multiplicity” is expressed like an Attribute “multiplicity” (see Table 3-2 on page 3-6), except that when it is absent, the “multiplicity” defaults to “exactly one.”</p> <p>The return Parameter (if it exists) should be the first contained Parameter of the Operation. If this entry is absent or says “none,” the Operation does not have a return Parameter.</p>
<i>isQuery:</i>	This optional entry defines the Operation’s “isQuery” flag. If it is absent, the “isQuery” flag has the value false.
<i>scope:</i>	This optional entry defines the Operation’s “scope.” If it is absent, the Operation has a “scope” of “instance_level.”
<i>parameters:</i>	This entry defines the Operation’s non-return Parameter list in the order that they appear in the Operation’s signature. The “name,” “direction,” “type,” and “multiplicity” are defined for each Parameter. If the “multiplicity” is not explicitly specified, it defaults to “exactly one.” If the entry simply says “none,” the Operation has no non-return Parameters.
<i>exceptions:</i>	This optional entry defines the list of Exceptions that this Operation may raise in the order that they appear in the Operation’s signature. If it is absent, the Operation raises no Exceptions.
<i>operation semantics:</i>	This optional entry simply gives a cross reference to the OCL defining the Operation’s semantics. Note that the MOF Model does not provide a standard way of representing an Operation’s semantic specification, and it is not included in the normative XMI serialization of the MOF Model.

3.2.1.7 Constraints

The “Constraints” heading lists the Constraints that are attached to this Class in the MOF Model. The OCL for the Constraints may be found in Section 3.9.3, “The MOF Model Constraints,” on page 3-107. Each listed Constraint “constrains” the Class, and is also contained by it.

3.2.1.8 IDL

The “IDL” heading shows an excerpt of the MOF Model IDL that corresponds to this Class. The excerpts, which are part of the “Model” module given in the “MOF to IDL Summary” appendix, consist of a Class proxy interface and an Instance interface. For information on these interfaces, refer to the “MOF to IDL Mapping” chapter.

3.2.2 Associations

The Associations in the MOF Model are defined in Section 3.5, “MOF Model Associations,” on page 3-77.

Associations describe relationships between instances of Classes. In short, an Association relates two Classes (or relates one Class to itself) to define a “link set” that contains two-ended “links” between instances of the Classes. The properties of an Association rest mostly in its two AssociationEnds. Refer to Section 4.7, “Semantics of Associations,” on page 4-12 for a more detailed explanation.

3.2.2.1 Association Heading

Each Association in the MOF Model is introduced by a second level section heading in Section 3.5, “MOF Model Associations,” on page 3-77. The heading defines the standard ModelElement name for the Association. The Classes name on the heading line can be followed by the word “*derived*.” For example, the following:

Chapter 3.5.4 Exposes *derived*

introduces an Association called “Exposes” and defines its Chapter “isDerived” flag to be true. If the word “*derived*” is not present, the Association’s “isDerived” flag is false.

The paragraph or paragraphs following an Association heading give a description of the Association, its purpose, and its meaning.

3.2.2.2 Ends

The “Ends” heading defines the two AssociationEnds for an Association in the MOF Model. The two AssociationEnds are defined by giving their “name” values and defining the remaining characteristics in tabular form.

Every AssociationEnd in the MOF Model has both “isNavigable” and “isChangeable” set to true. The remaining characteristics of AssociationEnds are defined using notation described in Table 3-4.

Table 3-4 AssociationEnds Characteristics

Entry	Description
<i>class:</i>	This entry specifies the Class whose instances are linked at this end of the Association. This is represented by the AssociationEnd’s “name” attribute.
<i>multiplicity:</i>	This entry defines the AssociationEnd’s “multiplicity” attribute. This is expressed in the same way as References (i.e., uniqueness is implicit - see Section 3.2.1.5, “References,” on page 3-5). Note the following: <ul style="list-style-type: none"> the OCL constraints on MultiplicityType and AssociationEnd mean that the “isUnique” field must be “false” if the “upper” bound is 1 and “true” otherwise, and the “multiplicity” settings for an AssociationEnd and its corresponding Reference(s) must be the same.
<i>aggregation:</i>	This optional entry defines the AssociationEnd’s “aggregation” attribute as one of “composite,” “shared,” or “none” (see Section 4.7.5, “Association Aggregation,” on page 4-16). If the entry is absent, the AssociationEnd’s “aggregation” attribute takes the value “none.”

3.2.2.3 Derivation

The “Derivation” heading defines how a derived Association should be computed. It may include marker for an OCL rule defined in Section 3.9.4, “Semantic specifications for some Operations, derived Attributes and Derived Associations,” on page 3-123.

3.2.2.4 IDL

The “IDL” heading shows an excerpt of the MOF Model IDL that corresponds to this Association. These excerpts, which are part of the “Model” module given in the MOF IDL Summary appendix, consist of an Association interface and related IDL data types. For more information, refer to the MOF to IDL Mapping chapter.

3.2.3 DataTypes

The DataTypes that form part of the MOF Model are described in Section 3.6, “MOF Model Data Types,” on page 3-93.

All DataTypes in the MOF Model have “visibility” of “public_vis.” The settings of the dummy attributes are “isAbstract” - false, “isRoot” - true, and “isLeaf” - true.

The remaining characteristics are the Exception’s

- “name” - given in the section heading,
- Container - given by the “Container” heading, and
- “typeCode” that can be determined from the declaration under the “IDL” heading.

If the “Container” heading is absent, the DataType is contained by the Model Package.

3.2.4 *Exceptions*

The Exceptions that form part of the MOF Model are described in Section 3.7, “MOF Model Exceptions,” on page 3-99.

All Exceptions in the MOF Model have “visibility” of “public_vis” and “scope” of “classifier_level.”

The remaining characteristics are the Exception’s

- “name” - given in the section heading, and
- Parameters and Container, which are given in the corresponding headings.

If the Container heading is absent, the Exception is contained by the Model Package.

3.2.5 *Constants*

The Constants that form part of the MOF Model are described in Section 3.9, “MOF Model Constraints and other M2 level semantics,” on page 3-102.

The characteristics of a Constant are its

- “name” - given in the section heading, and
- Container - given under the “Container” heading, and
- “type” and “value” that can be determined from the IDL.

If the “Container” heading is absent, the DataType is contained by the Model Package.

3.2.6 *Constraints*

The Constraints that form part of the MOF Model are described in Section 3.9, “MOF Model Constraints and other M2 level semantics,” on page 3-102. The notation used for describing the constraints is described in Section 3.9.1.1, “Notation for MOF Model Constraints,” on page 3-103.

3.2.7 *UML Diagrams*

At various points in this chapter, UML class diagrams are used to describe aspects of the MOF Model. To understand these diagrams, the reader should mentally map from UML modeling concepts to the equivalent MOF meta-modelling constructs.

There is one point in which this document's use of UML notation requires explaining. In standard UML notation, an arrowhead on an Association line indicates that the Association is navigable in the direction indicated. Absence of an arrowhead can mean either that the Association is navigable or that it is navigable in both directions, depending on the context.

As was explained in Section 3.2.1.5, "References," on page 3-5, the MOF models navigable Associations in a different way. Thus in this document, an arrowhead on one end of an Association means that a Reference exists on the Class at the opposite end that allows navigation in the indicated direction. If there are no arrowheads, there are References on the Classes at *both* ends of the Association.

3.3 *The Structure of the MOF Model*

This section gives an overview of the structure of the MOF Model.

3.3.1 *The MOF Model Package*

The MOF Model, as it is currently defined, consists of a single non-nested Package called "Model."

The class diagram in Figure 3-1 on page 3-12 shows the Classes and Associations of the "Model" Package. To aid readability, Class features, DataTypes, and other details have been omitted from the diagram. These details are all specified in later sections of this chapter.

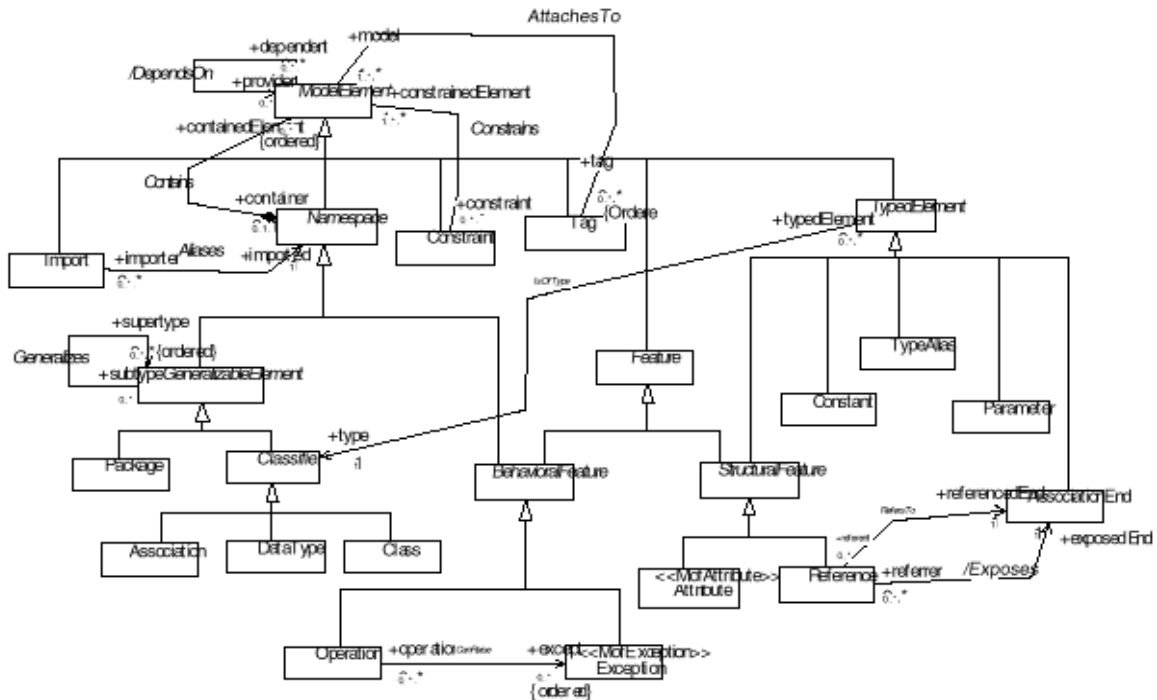


Figure 3-1 The MOF Model Package

3.3.2 The MOF Model Service IDL

The “Model” Package is used to generate the CORBA IDL for the OMG MOF Model service using the MOF to IDL Mapping defined in the “MOF Abstract Mapping” chapter through the “Reflective Module” chapter. Relevant fragments of the resulting IDL is embedded in the Class, Association, DataType, and Exception descriptions in sections Section 3.4, “MOF Model Classes,” on page 3-15 through Section 3.7, “MOF Model Exceptions,” on page 3-99. The complete IDL, along with the dependent “Reflective” IDL may be found in the MOF IDL Summary appendix.

The IDL for the MOF Model service requires a “prefix” of “org.omg.mof.” To this end, the “Model” Package is defined to have an “idl_prefix” Tag with value “org.omg.mof.”

3.3.3 The MOF Model Structure

The core structure of the MOF Model is shown in the class diagram in Figure 3-2. This diagram shows the key abstract Classes in the MOF Model and the key Associations between them.

3.3.3.1 Key Abstract Classes

The key abstract Classes in the MOF Model are as follows:

- ModelElement - this is the common base Class of all M3-level Classes in the MOF Model. Every ModelElement has a “name.”
- Namespace - this is the base Class for all M3-level Classes that need to act as containers in the MOF Model.
- GeneralizableElement - this is the base Class for all M3-level Classes that support “generalization” (i.e., inheritance).
- TypedElement - this is the base Class for M3-level Classes such as Attribute, Parameter, and Constant whose definition requires a type specification.
- Classifier - this is the base Class for all M3-level Classes that (notionally) define types. Examples of Classifier include Class and DataType.

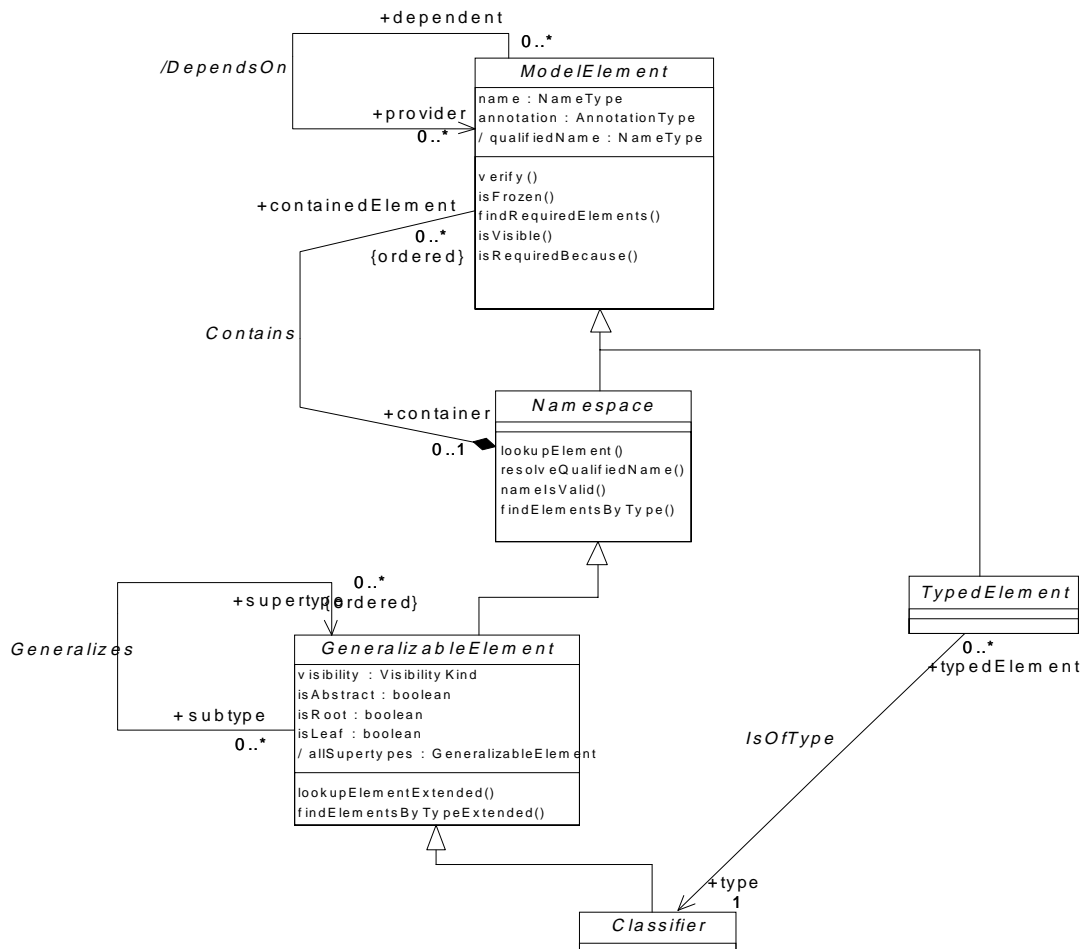


Figure 3-2 The Key Abstractions of the MOF Model

3.3.3.2 Key Associations

The key Associations in the MOF Model are as follows:

- Contains - this Association relates a ModelElement to the Namespace that contains it (see Section 3.3.4, “The MOF Model Containment Hierarchy,” on page 3-14).
- Generalizes - this Association relates a GeneralizableElement to its ancestors (i.e., supertypes) and children (i.e., subtypes) in a model element inheritance graph. Note that a GeneralizableElement may not know about all of its subtypes.
- IsOfType - this Association relates a TypedElement to the Classifier that defines its type.
- DependsOn - this derived Association relates a ModelElement to others that its definition depends on. (It is derived from Contains, Generalizes, IsOfType, and other Associations not shown here.)

3.3.4 The MOF Model Containment Hierarchy

The most important relationship in the MOF Model is the Contains Association. Containment is a utility Association that is used to relate (for example) Classes to their Operations and Attributes, Operations to their Parameters and so on. While the class diagram shows that only ModelElement objects which are subtypes of Namespace can contain any other ModelElements, the MOF Model restricts the legal containments to eliminate various non-sensical and problematical cases.

Table 3-5 defines the legal ModelElement containments in matrix form. The row headings list all non-abstract subtypes of Namespace (i.e., possible containers) and the column headings list all non-abstract ModelElements (i.e., possible contained elements). For each combination of container and contained, the corresponding table cell is either “Y” showing that containment is legal or “N” showing that it is not.

Table 3-5 The ModelElement Containment Matrix

	Package	Class	Data Type	Association	Attribute	Reference	Operation	Exception	Parameter	Association End	Constraint	Constant	Type Alias	Tag
Package	Y	Y	Y	N	N	N	N	Y	N	N	Y	Y	N	Y
Class	N	Y	Y	N	Y	Y	Y	Y	N	N	Y	Y	N	Y
Data Type	N	N	N	N	N	N	N	N	N	N	Y	N	Y	Y
Association	N	N	N	N	N	N	N	N	N	Y	Y	N	N	Y
Operation	N	N	N	N	N	N	N	N	Y	N	Y	N	N	Y
Exception	N	N	N	N	N	N	N	N	Y	N	N	N	N	Y

3.4 MOF Model Classes

3.4.1 ModelElement Class

ModelElement

abstract

ModelElement classifies the elementary, atomic constructs of models. ModelElement is the root Class within the MOF Model.

SuperClasses

None. (While the CORBA IDL for ModelElement inherits from Reflective::RefObject, this is *not* generalization in the MOF Model sense. Rather it is an artifact of the IDL mapping.)

Attributes

name

Provides a meta-modeler supplied name that uniquely identifies the ModelElement in the context of the ModelElement's containing Namespace. When choosing a ModelElement's name, the meta-modeler should consider the rules for translating names into identifiers in the relevant mappings (e.g., Section 5.7.1, "Generated IDL Identifiers," on page 5-39).
--

<i>type:</i>	NameType
--------------	----------

<i>multiplicity:</i>	exactly one
----------------------	-------------

annotation

Provides an informal description of the ModelElement..
--

<i>type:</i>	AnnotationType
--------------	----------------

<i>multiplicity:</i>	exactly one
----------------------	-------------

qualifiedName

Provides a unique name for the ModelElement within the context of its outermost containing Package. The qualifiedName is a list of NameType values consisting of the names of the ModelElement, its container, its container's container and so on until a non-contained element is reached. The first member of the list is the name of the non-contained element.

<i>type:</i>	NameType
--------------	----------

<i>multiplicity:</i>	one or more; ordered; non-unique
<i>changeable:</i>	no
<i>derived from:</i>	[S-12]

References

container

Identifies the Namespace that contains the ModelElement. Since the Contains Association is a Composite Association, any ModelElement can have at most one container, and the containment graph is strictly tree shaped.	
<i>type:</i>	Namespace
<i>defined by:</i>	Contains::container
<i>multiplicity:</i>	zero or one
<i>inverse:</i>	ModelElement::contents

requiredElements

Identifies the ModelElements on whose definition the definition of this ModelElement depends. For a definition of dependency, see Section 3.5.9, “DependsOn,” on page 3-89.	
<i>type:</i>	ModelElement
<i>defined by:</i>	DependsOn::provider
<i>multiplicity:</i>	zero or more

constraints

Identifies the set of Constraints that apply to the ModelElement. A Constraint applies to all instances of the ModelElement and its sub-Classes.	
<i>type:</i>	Constraint
<i>multiplicity:</i>	zero or more
<i>inverse:</i>	Constraint::constrainedElements.
<i>defined by:</i>	Constrains::provider

Operations

verify

Each ModelElement is capable of checking its own correctness, as defined by the inherent properties of meta-models described in this specification, and constraints that hold over the ModelElement. The client of the operation specifies whether the operation should propagate to any ModelElements that this ModelElement might contain (if it is capable of containing elements), or whether it should return after only checking itself. The verify operation checks inherent constraints on the object and its attributes plus any constraints contained by the object. The operation returns valid if all verification checks passed; otherwise, it returns invalid. A parameter returns representations of any constraint violations detected. If the operation returns invalid, this parameter must not be empty. When the depth argument is deep, and this element (and, by definition, all its dependent elements) are published, the operation returns published.

<i>return type:</i>	VerifyResultKind
<i>parameters:</i>	depth: in DepthKind violations: out ViolationType (multiplicity: zero or more; unique)
<i>isQuery:</i>	yes

isFrozen

Reports the freeze status of a ModelElement. A ModelElement, at any particular time, is either frozen or not frozen. All ModelElements of a published model are permanently frozen.

<i>return type:</i>	boolean
<i>isQuery:</i>	yes

findRequiredElements

Supports selecting a subset of the ModelElements that this one depends on, based on their dependency categories. The “kinds” argument gives a bag of DependencyKinds that are of interest to the caller. String constants for the standard dependency categories are given in Section 3.8, “MOF Model Constants,” on page 3-95 and their meanings are defined in Section 3.5.9, “DependsOn,” on page 3-84. In this context, the AllDep pseudo-category (i.e., “all”) is equivalent to passing all of the standard categories in the “kinds” argument, and the IndirectDep pseudo-category (i.e., “indirect”) is ignored.

If the “recursive” argument is “false”, the operation return the direct dependents only. If it is “true” all dependents in the transitive closure of the specified “kinds” are returned.

<i>return type:</i>	ModelElement (multiplicity: zero or more; unordered, unique)
<i>isQuery:</i>	yes
<i>parameters:</i>	kinds: in DependencyKind (multiplicity: one or more; unordered; unique) recursive: in boolean
<i>operation semantics</i>	[S-4]

isVisible

Returns true. This operation is reserved for future use when the MOF visibility rules have stabilised. Then it will determine whether the supplied otherElement is visible to this ModelElement.

<i>return type:</i>	boolean
<i>isQuery:</i>	yes
<i>parameters:</i>	otherElement: in ModelElement
<i>operation semantics</i>	[S-3]

isRequiredBecause

This operation performs two functions:

- It checks whether this ModelElement directly or indirectly depends on the ModelElement given by “otherElement”. If it does, the operation’s result is “true;” otherwise, it is “false.”
- If a dependency exists (i.e., the result is “true”), the operation returns a DependencyKind in “reason” that categorizes the dependency. String constants for the DependencyKind categories are given in Section 3.8, “MOF Model Constants,” on page 3-95 and their meanings are defined in Section 3.5.9, “DependsOn,” on page 3-84. If the dependency is indirect, IndirectDep is returned. If there are multiple dependencies, any category that applies may be returned in “reason.” If no dependencies exist, an empty string is returned in “reason.”

<i>return type:</i>	boolean
<i>isQuery:</i>	yes
<i>parameters:</i>	otherElement: in ModelElement reason: out DependencyKind
<i>operation semantics</i>	[S-5]

Constraints

A ModelElement that is not a Package must have a container. [C-1]

The attribute values of a ModelElement which is frozen cannot be changed. [C-2]

A frozen ModelElement which is in a frozen Namespace can only be deleted, by deleting the Namespace. [C-3]

The link sets that express dependencies of a frozen Element on other Elements cannot be explicitly changed. [C-4]

IDL

```
interface ModelElementClass : Reflective::RefObject {
  readonly attribute ModelElementUList all_of_type_model_element;
  const string MUST_BE_CONTAINED_UNLESS_PACKAGE =
    "org.omg.mof:constraint.model.model_element.must_be_contained_unless_package";
  const string FROZEN_ATTRIBUTES_CANNOT_BE_CHANGED =
    "org.omg.mof:constraint.model.model_element.frozen_attributes_cannot_be_changed";
  const string FROZEN_ELEMENTS_CANNOT_BE_DELETED =
    "org.omg.mof:constraint.model.model_element.frozen_elements_cannot_be_deleted";
  const string FROZEN_DEPENDENCIES_CANNOT_BE_CHANGED =
    "org.omg.mof:constraint.model.model_element.frozen_dependencies_cannot_be_changed";
  typedef string DependencyKind;
  typedef sequence <::Model::ModelElementClass::DependencyKind>DependencyKindSet;
  const DependencyKind CONTAINER_DEP = "container";
  const DependencyKind CONTENTS_DEP = "contents";
  const DependencyKind SIGNATURE_DEP = "signature";
  const DependencyKind CONSTRAINT_DEP = "constraint";
  const DependencyKind CONSTRAINED_ELEMENTS_DEP = "constrained elements";
```

```

const DependencyKind SPECIALIZATION_DEP = "specialization";
const DependencyKind IMPORT_DEP = "import";
const DependencyKind TYPE_DEFINITION_DEP = "type definition";
const DependencyKind REFERENCED_ENDS_DEP = "referenced ends";
const DependencyKind TAGGED_ELEMENTS_DEP = "tagged elements";
const DependencyKind INDIRECT_DEP = "indirect";
const DependencyKind ALL_DEP = "all";
enum VerifyResultKind {valid, published, invalid};
enum DepthKind {shallow, deep};
struct ViolationType {
string error_kind;
Reflective::RefObject element_in_error;
Reflective::NamedValueList values_in_error;
string error_description;
};
typedef sequence <::Model::ModelElementClass::ViolationType> ViolationTypeSet;
}; // end of interface ModelElementClass

```

```

interface ModelElement : ModelElementClass {
NameType name ()
raises (Reflective::MofError);
void set_name (in NameType new_value)
raises (Reflective::MofError);
NameTypeList qualified_name ()
raises (Reflective::MofError);
AnnotationType annotation ()
raises (Reflective::MofError);
void set_annotation (in AnnotationType new_value)
raises (Reflective::MofError);
ModelElementSet required_elements ()
raises (Reflective::MofError);
ModelElementSet find_required_elements (
in ModelElementClass::DependencyKindSet kinds,
in boolean recursive)
raises (Reflective::MofError);
boolean is_required_because (in ModelElement other_element,
out ModelElementClass::DependencyKind reason)
raises (Reflective::MofError);
Namespace container ()
raises (Reflective::NotSet, Reflective::MofError);
void set_container (in Namespace new_value)
raises (Reflective::MofError);
void unset_container ()
raises (Reflective::MofError);
ConstraintSet constraints ()
raises (Reflective::MofError);
void set_constraints (in ConstraintSet new_value)
raises (Reflective::MofError);
void add_constraints (in Constraint new_element)
raises (Reflective::MofError);
void modify_constraints (in Constraint old_element, in Constraint new_element)
raises (Reflective::MofError);
void remove_constraints (in Constraint old_element)
raises (Reflective::NotFound, Reflective::MofError);
ModelElementClass::VerifyResultKind verify (

```

```

        in ModelElementClass::DepthKind depth,
        out ModelElementClass::ViolationTypeSet problems)
    raises (Reflective::MofError);
boolean is_frozen ()
    raises (Reflective::MofError);
boolean is_visible (in ModelElement other_element)
    raises (Reflective::MofError);
};

```

3.4.2 Namespace Class

Namespace

abstract

The Namespace Class classifies and characterizes ModelElements that can contain other ModelElements. Along with containing the ModelElements, a Namespace defines a namespace, the allowable set of names and the naming constraints, for these elements.

Subclasses of the Namespace Class have mechanisms for effectively extending their namespace, without actually containing additional ModelElements. Thus Namespace can be viewed in terms of its two roles, as a container and as a namespace mechanism. Because only subclasses extend the namespace, the namespace and contents are coincident in the definition of the Namespace Class. Each Namespace has four collections (the latter three derivable) that are used in the MOF Model's Constraints. These collections are:

- The contents (also called the direct contents), which are defined by the contents reference.
- All contents, the transitive closure on the contents reference.
- The extended namespace (the contents plus elements included by extension), which Namespace subclasses accomplish through generalization and importation.
- The extended contents (the transitive closure on the contents reference applied to the extended namespace).

The definitions of these collections may be found in Section 3.9.5, "OCL Helper functions," on page 3-129.

SuperClasses

ModelElement

References

contents

Identifies the set of ModelElements that a Namespace contains.	
<i>class:</i>	ModelElement
<i>defined by:</i>	Contains::containedElement
<i>multiplicity:</i>	zero or more; ordered
<i>inverse:</i>	ModelElement::container

Operations

lookupElement

Searches for an element contained by this Namespace whose name is precisely equal (as a string) to the supplied name. The operation either returns a ModelElement that satisfies the above, or raises the NameNotFound exception.	
<i>return type:</i>	ModelElement
<i>isQuery:</i>	yes
<i>parameters:</i>	name : in NameType
<i>exceptions:</i>	NameNotFound
<i>operation semantics:</i>	[S-6]

resolveQualifiedName

Searches for a ModelElement contained within this Namespace that is identified by the supplied qualifiedName. The qualifiedName is interpreted as a “path” starting from this Namespace.	
<i>return type:</i>	ModelElement (exactly one). If no element is found, an exception is raised.
<i>isQuery:</i>	yes
<i>parameters:</i>	qualifiedName : in NameType (multiplicity one or more; ordered; not unique)
<i>exceptions:</i>	NameNotResolved
<i>operation semantics:</i>	[S-7]

nameIsValid

Determines whether the proposedName can be used as the name for a new member ModelElement in this Namespace. Specifically, it checks that the Namespace uniqueness rules would still be satisfied after adding such a name.	
<i>return type:</i>	boolean
<i>isQuery:</i>	yes
<i>parameters:</i>	proposedName : in NameType
<i>operation semantics:</i>	[S-8]

findElementsByType

Returns all the ModelElements identified by the contents reference defined for this Namespace that are of the Class supplied. The returned list of ModelElements is a subset of the ModelElements contained by this Namespace. This operation can either return only those ModelElements that exactly match the specified class or those ModelElements that are instances of the specified class and one or more of its subclasses.	
Because ModelElement is an abstract class, invoking this operation with the ofType argument specified as ModelElement and the includeSubtypes argument set to false returns an empty list. Because ModelElement is the base type for all instances that can be contained by a Namespace, invoking the operation with the ofType argument specified as ModelElement, and includeSubtypes set to true, returns all the contained elements of the Namespace.	
<i>return type:</i>	ModelElement (multiplicity zero or more; ordered; unique)
<i>isQuery:</i>	yes
<i>parameters:</i>	ofType : in Class includeSubtypes : in boolean
<i>operation semantics:</i>	[S-9]

Constraints

The names of the contents of a Namespace must not collide. [C-5]

IDL

```
interface NamespaceClass : ModelElementClass {
  readonly attribute NamespaceUList all_of_type_namespace;
  const string CONTENT_NAMES_MUST_NOT_COLLIDE =
    "org.omg.mof.constraint.model.namespace.content_names_must_not_collide";

  exception NameNotFound {
```

```

    NameType name;
};
exception NameNotResolved {
    string explanation;
    NameTypeList rest_of_name;
};
}; // end of interface NamespaceClass

interface Namespace : NamespaceClass, ModelElement {
    ModelElementUList contents ()
        raises (Reflective::MofError);
    void set_contents (in ModelElementUList new_value)
        raises (Reflective::MofError);
    void add_contents (in ModelElement new_element)
        raises (Reflective::MofError);
    void add_contents_before (in ModelElement new_element,
        in ModelElement before_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_contents (in ModelElement old_element,
        in ModelElement new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_contents (in ModelElement old_element)
        raises (Reflective::NotFound, Reflective::MofError);

    ModelElement lookup_element (in NameType name)
        raises (NamespaceClass::NameNotFound, Reflective::MofError);
    ModelElement resolve_qualified_name (in NameTypeList qualified_name)
        raises (NamespaceClass::NameNotResolved, Reflective::MofError);
    ModelElementUList find_elements_by_type (in Class of_type,
        in boolean include_subtypes)
        raises (Reflective::MofError);
    boolean name_is_valid (in NameType proposed_name)
        raises (Reflective::MofError);
};

```

3.4.3 GeneralizableElement Class

GeneralizableElement

abstract

The GeneralizableElement Class classifies and characterizes ModelElements that can be generalized through supertyping and specialized through subtyping. A GeneralizableElement inherits the features of each of its supertypes, the features of the supertypes of the immediate supertypes, and so on (in other words all the features of the transitive closure of all the supertypes of the GeneralizableElement).

When a GeneralizableElement inherits a feature, that feature name effectively becomes part of the namespace for the GeneralizableElement and the feature is considered part of the extended namespace of the Namespace. Therefore, a GeneralizableElement cannot have a superclass if it causes an inherited feature to have a namespace collision with its own features - see Constraint [C-8].

To the degree that a GeneralizableElement is defined by its features, the superclass/subclass association defines substitutability. Any instance of a GeneralizableElement can be supplied wherever an instance of a superclass of that GeneralizableElement is expected.

SuperClasses

Namespace

Attributes

visibility

In the future, this Attribute will be used to limit the ability of ModelElements outside of this GeneralizableElement's container to depend on it; see Section 3.6.6, "VisibilityKind," on page 3-95. The rules of visibility of MOF ModelElements are not currently specified.	
<i>type:</i>	VisibilityKind
<i>multiplicity:</i>	exactly one

isAbstract

Indicates whether the GeneralizableElement is expected to have instances. When isAbstract is true, any instance that is represented or classified by this GeneralizableElement is additionally an instance of some specialization of this GeneralizableElement. No operation that supports creation of instances of this GeneralizableElement should be available.	
<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

isRoot

Specifies whether the GeneralizableElement may have supertypes. True indicates that it may not have supertypes, false indicates that it may have supertypes (whether or not it actually has any)	
<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

isLeaf

Specifies whether the GeneralizableElement may be a supertype of another Generalizable Element. True indicates that it may not be a supertype, false indicates that it may be a supertype (whether or not it actually is).	
<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

References***supertypes***

Identifies the set of superclasses for a GeneralizableElement. Note that a GeneralizableElement does not have a reference to its subclasses.	
<i>class:</i>	GeneralizableElement
<i>defined by:</i>	Generalizes::supertype
<i>multiplicity:</i>	zero or more; ordered

Operations***allSupertypes***

Returns a list of direct and indirect supertypes of this GeneralizableElement. A direct supertype is a GeneralizableElement that directly generalizes this one. An indirect supertype is defined (recursively) as a supertype of some other direct or indirect supertype of the GeneralizableElement. The order of the list elements is determined by a depth-first traversal of the supertypes with duplicate elements removed.	
<i>return type:</i>	GeneralizableElement (multiplicity zero or more, ordered, unique)
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>operation semantics:</i>	[S-1]

lookupElementExtended

Returns an element whose name matches the supplied “name.” Like the “lookupElement” operation on Namespace, this operation searches the contents of the GeneralizableElement. In addition, it tries to match the name in the contents of all direct and indirect supertypes of the GeneralizableElement. For Packages, a subclass of GeneralizableElement, the operation can also match a Namespace associated with an Import objects. NameNotFound is raised if no element matches the name.	
<i>return type:</i>	ModelElement (multiplicity exactly one)
<i>isQuery:</i>	yes
<i>parameters:</i>	name : in NameType
<i>exceptions</i>	NameNotFound
<i>operation semantics:</i>	[S-10]

findElementsByTypeExtended

Provides an extension of the findElementsByType defined for Namespace so that contained elements of all superclasses (direct and indirect) of the GeneralizableElement are included in the search. The order of the returned elements is determined by the order of the elements contained in the GeneralizableElements and a depth-first traversal of the superclasses.	
Subclasses can include a larger overall area for the lookup. Package, a subclass of GeneralizableElement, also considers the elements brought into this Namespace through the use of Import.	
<i>return type:</i>	ModelElement (multiplicity zero or more; ordered; unique)
<i>isQuery:</i>	yes
<i>parameters:</i>	ofType : in Class includeSubtypes : in boolean
<i>operation semantics:</i>	[S-11]

Constraints

A Generalizable Element cannot be its own direct or indirect supertype. [C-6]

A supertypes of a GeneralizableElement must be of the same kind as the GeneralizableElement itself. [C-7]

The names of the contents of a GeneralizableElement should not collide with the names of the contents of any direct or indirect supertype. [C-8]

Multiple inheritance must obey the “Diamond Rule.” [C-9]

If a Generalizable Element is marked as a “root,” it cannot have any supertypes. [C-10]

A GeneralizableElement’s immediate supertypes must all be visible to it. [C-11]

A GeneralizableElement cannot inherit from a GeneralizableElement defined as a “leaf.” [C-12]

IDL

```
interface GeneralizableElementClass : NamespaceClass {
  readonly attribute GeneralizableElementUList
  all_of_type_generalizable_element;
  const string SUPERTYPE_MUST_NOT_BE_SELF =
    "org.omg.mof:constraint.model.generalizable_element.supertype_must_not_be_self";
  const string SUPERTYPE_KIND_MUST_BE_SAME =
    "org.omg.mof:constraint.model.generalizable_element.supertype_kind_must_be_same";
  const string CONTENTS_MUST_NOT_COLLIDE_WITH_SUPERTYPES
    "org.omg.mof:constraint.model.generalizable_element"
    ".contents_must_not_collide_with_supertypes";
  const string DIAMOND_RULE_MUST_BE_OBEYED =
    "org.omg.mof:constraint.model.generalizable_element.diamond_rule_must_be_obeyed";
  const string NO_SUPERTYPES_ALLOWED_FOR_ROOT =
    "org.omg.mof:constraint.model.generalizable_element.no_supertypes_allowed_for_root";
  const string SUPERTYPES_MUST_BE_VISIBLE =
    "org.omg.mof:constraint.model.generalizable_element.supertypes_must_be_visible";
  const string NO_SUBTYPES_ALLOWED_FOR_LEAF =
    "org.omg.mof:constraint.model.generalizable_element.no_subtypes_allowed_for_leaf";
}; // end of interface GeneralizableElementClass

interface GeneralizableElement : GeneralizableElementClass, Namespace {
  boolean is_root ()
    raises (Reflective::MofError);
  void set_is_root (in boolean new_value)
    raises (Reflective::MofError);
  boolean is_leaf ()
    raises (Reflective::MofError);
  void set_is_leaf (in boolean new_value)
    raises (Reflective::MofError);
  boolean is_abstract ()
    raises (Reflective::MofError);
  void set_is_abstract (in boolean new_value)
    raises (Reflective::MofError);
  VisibilityKind visibility ()
    raises (Reflective::MofError);
  void set_visibility (in VisibilityKind new_value)
    raises (Reflective::MofError);
  GeneralizableElementUList supertypes ()
    raises (Reflective::MofError);
  void set_supertypes (in GeneralizableElementUList new_value)
    raises (Reflective::MofError);
  void add_supertypes (in GeneralizableElement new_element)
    raises (Reflective::MofError);
  void add_supertypes_before (in GeneralizableElement new_element,
```

```

        in GeneralizableElement before_element)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_supertypes (in GeneralizableElement old_element,
    in GeneralizableElement new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_supertypes (in GeneralizableElement old_element)
    raises (Reflective::NotFound, Reflective::MofError);
GeneralizableElementUList all_supertypes ()
    raises (Reflective::MofError);
ModelElement lookup_element_extended (in NameType name)
    raises (NameNotFound, Reflective::MofError);
ModelElementUList find_elements_by_type_extended (in Class of_type,
    in boolean include_subtypes)
    raises (Reflective::MofError);
};

```

3.4.4 TypedElement

TypedElement

abstract

The TypedElement type is an abstraction of ModelElements that require a type as part of their definition. A TypedElement does not itself define a type, but is associated with a Classifier.

SuperClasses

ModelElement

References

type

Provides the representation of the type supporting the TypedElement through this reference.	
<i>class:</i>	Classifier
<i>defined by:</i>	IsOfType::type
<i>multiplicity:</i>	exactly one

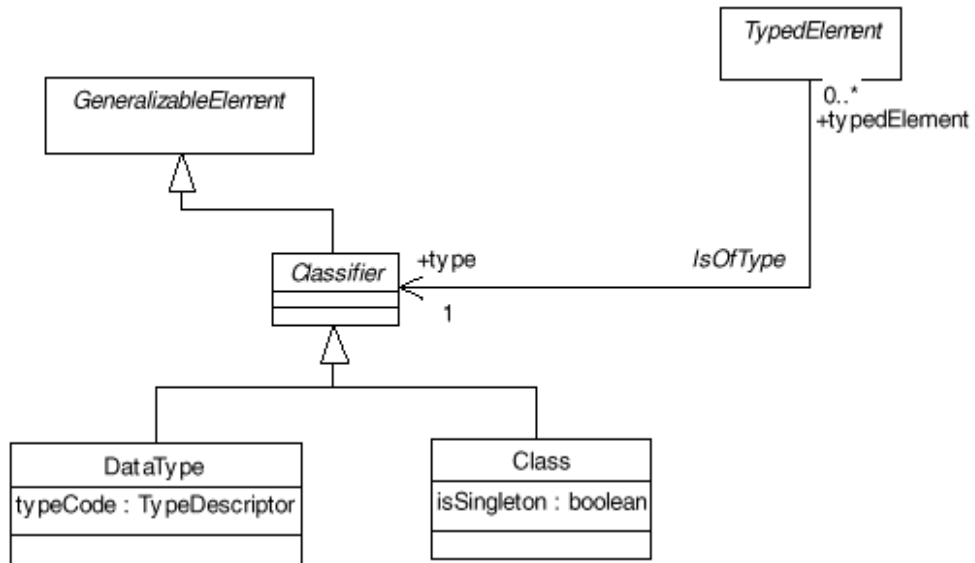


Figure 3-3 MOF Model Classifiers

Constraints

An Association cannot be the type of a TypedElement. [C-13]

A TypedElement can only have a type that is visible to it. [C-14]

IDL

```

interface TypedElementClass : ModelElementClass {
    // get all typed_element including subtypes of typed_element
    readonly attribute TypedElementUList all_of_type_typed_element;
    const string ASSOCIATIONS_CANNOT_BE_TYPES =
        "org.omg.mof.constraint.model.typed_element.associations_cannot_be_types";
    const string TYPE_MUST_BE_VISIBLE =
        "org.omg.mof.constraint.model.typed_element.type_must_be_visible";
}; // end of interface TypedElementClass

interface TypedElement : TypedElementClass, ModelElement {
    Classifier type ()
        raises (Reflective::MofError);
    void set_type (in Classifier new_value)
        raises (Reflective::MofError);
};
  
```

3.4.5 Classifier

Classifier

abstract

A classifier provides a classification of instances through a set of Features it contains.

SuperClasses

GeneralizableElement

IDL

```
interface ClassifierClass : GeneralizableElementClass {
    readonly attribute ClassifierULList all_of_type_classifier;
}; // end of interface ClassifierClass
```

```
interface Classifier : ClassifierClass, GeneralizableElement { };
```

3.4.6 Class

A Class defines a classification over a set of object instances by defining the behavior they exhibit. This behavior is represented through operations, attributes, references, participation in associations, nested classes, constants, and constraints. Although the same or similar elements are used in other environments for representing Classes and their implementations, in the MOF these elements specify the class characteristics in an implementation-independent manner. For instance, defining a Class as having an attribute does not require the implementation of the Class (the software which provides the conformant behavior) to have an attribute, to hold the attribute value, etc. The implementation simply must insure that the behavior conforms to the definition of the attribute. The use of all the additional elements beyond operations provides a much richer environment for defining Class behavior. Likewise, this construct is not an interface. Its expressibility goes beyond interface representations. The MOF's IDL translation capabilities map a single Class onto two interfaces. It would be possible to define transformations to alternate interface representations, such as Java's interfaces.

There are, and probably always will be, divergent views among industry leaders on the definition of the concepts of Class, Type, and Interface. As a domain-specific modeling environment, so long as the MOF is clear about the meaning of Class within the MOF, it should remain immune from such concerns.

SuperClasses

Classifier

Contained Elements

Class, DataType, Attribute, Reference, Operation, Exception, Constraint, Constant, Tag
- see [C-15].

Attributes

isSingleton

When isSingleton is true, at most one M1 level instance of this Class may exist within the M1-level extent of the Class.	
<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

Constraints

A Class may contain only Classes, DataTypes, Attributes, References, Operations, Exceptions, Constraints and Tags. [C-15]

A Class that is marked as abstract cannot also be marked as singleton. [C-16]

IDL

```
interface ClassClass : ClassifierClass {
  readonly attribute ClassUList all_of_type_class;
  readonly attribute ClassUList all_of_class_class;
  const string CLASS_CONTAINMENT_RULES =
    "org.omg.mof.constraint.model.class.class_containment_rules";
  const string ABSTRACT_CLASSES_CANNOT_BE_SINGLETON =
    "org.omg.mof.constraint.model.class.abstract_classes_cannot_be_singleton";

  Class create_class (
    /* from ModelElement */ in ::Model::NameType name,
    /* from ModelElement */ in ::Model::AnnotationType annotation,
    /* from GeneralizableElement */ in boolean is_root,
    /* from GeneralizableElement */ in boolean is_leaf,
    /* from GeneralizableElement */ in boolean is_abstract,
    /* from GeneralizableElement */ in ::Model::VisibilityKind visibility,
    /* from Class */ in boolean is_singleton)
    raises (Reflective::MofError);
}; // end of interface ClassClass

interface Class : ClassClass, Classifier {
  boolean is_singleton ()
    raises (Reflective::MofError);
  void set_is_singleton (in boolean new_value)
    raises (Reflective::MofError);
};
```


3.4.7 *DataType*

The *DataType* model element is primarily used to represent MOF data types and native types, as classified and described in Section 4.2, “MOF Values,” on page 4-2. The type information in a *DataType* is currently represented as a CORBA *TypeCode*, which has an encoding that is self-contained, transmissible, and relatively compact. *DataTypes* are also used to represent the type of a non-aggregate *Attribute* whose effective type is a *Class*.

The types represented by *DataTypes* fall into two groups:

1. A *DataType* that requires an IDL declaration must have a “name,” which maps to an IDL identifier.
2. A *DataType* that does not require an IDL declaration must have a “name” that starts with a “*” character. Such a *DataType* can be further classified as
 - a *DataType* for an anonymous CORBA data type (e.g., “boolean,” “char,” and so on),
 - a *DataType* that denotes a non-aggregate *Class* (i.e., one whose “typeCode” has kind of *tk_objRef* and that has a *TypeAlias* linkage to a *Class* - see Section 3.4.8, “*TypeAlias*,” on page 3-41), or
 - a *DataType* that denotes a named external CORBA interface or data type (e.g., a use of a CORBA type defined in the Interface Repository). External anonymous data types cannot be used.

In the last case, the IDL mapping relies on the *DataType*’s “typeCode” having a *repositoryId* that can be translated into a qualified name for the type.

Note – *TypeCodes* in a *DataType* are restricted to those supported by CORBA 2.2. (This restriction may be removed in a future revision of this document.) Furthermore, *TypeCodes* with kind of *tk_null*, *tk_void*, *tk_principal* and *tk_except* may not be used.

SuperClasses

Classifier

Contained Elements

TypeAlias, *Constraint*, *Tag*; see [C-17].

Attributes

typeCode

This attribute uses a CORBA TypeCode type to encode a DataType's type information. The TypeCode value should be such that the obvious 1-to-1 mapping between a TypeCode and IDL text applies. For example:

- <tk_boolean> should denote "boolean," and
- <tk_alias('Foo',<tk_boolean>)> should denote "typedef boolean Foo."

<i>type:</i>	TypeDescriptor
<i>multiplicity:</i>	exactly one

Constraints

A DataType may contain only TypeAliases, Constraints, and Tags. [C-17]

The typeCode of a DataType must denote a CORBA 2.2 compliant object type or data type. [C-18]

Inheritance / generalization is not applicable to DataTypes. [C-19]

A DataType cannot be abstract. [C-20]

IDL

```
interface DataTypeClass : ClassifierClass {
    readonly attribute DataTypeUList all_of_type_data_type;
    readonly attribute DataTypeUList all_of_class_data_type;
    const string DATA_TYPE_CONTAINMENT_RULES =
        "org.omg.mof:constraint.model.data_type.data_type_containment_rules";
    const string THIS_TYPECODE_NOT_SUPPORTED =
        "org.omg.mof:constraint.model.data_type.this_typecode_not_supported";
    const string DATA_TYPES_HAVE_NO_SUPERTYPES =
        "org.omg.mof:constraint.model.data_type.data_types_have_no_supertypes";
    const string DATA_TYPES_CANNOT_BE_ABSTRACT =
        "org.omg.mof:constraint.model.data_type.data_types_cannot_be_abstract";

    DataType create_data_type (
        /* from ModelElement */in ::Model::NameType name,
        /* from ModelElement */in ::Model::AnnotationType annotation,
        /* from GeneralizableElement */in boolean is_root,
        /* from GeneralizableElement */in boolean is_leaf,
        /* from GeneralizableElement */in boolean is_abstract,
        /* from GeneralizableElement */in ::Model::VisibilityKind visibility,
        /* from DataType */in TypeDescriptor type_code)
        raises (Reflective::MofError);
}; // end of interface DataTypeClass

interface DataType : DataTypeClass, Classifier {
    TypeDescriptor type_code ()
        raises (Reflective::MofError);
};
```

```

void set_type_code (in TypeDescriptor new_value)
  raises (Reflective::MofError);
};

```

3.4.8 TypeAlias

A TypeAlias is used to relate an embedded use of a type within a DataType to the Classifier that defines it. This is illustrated in Figure 3-4, which shows how a non-aggregate Class-valued Attribute would be represented in a/n MOF meta-model.

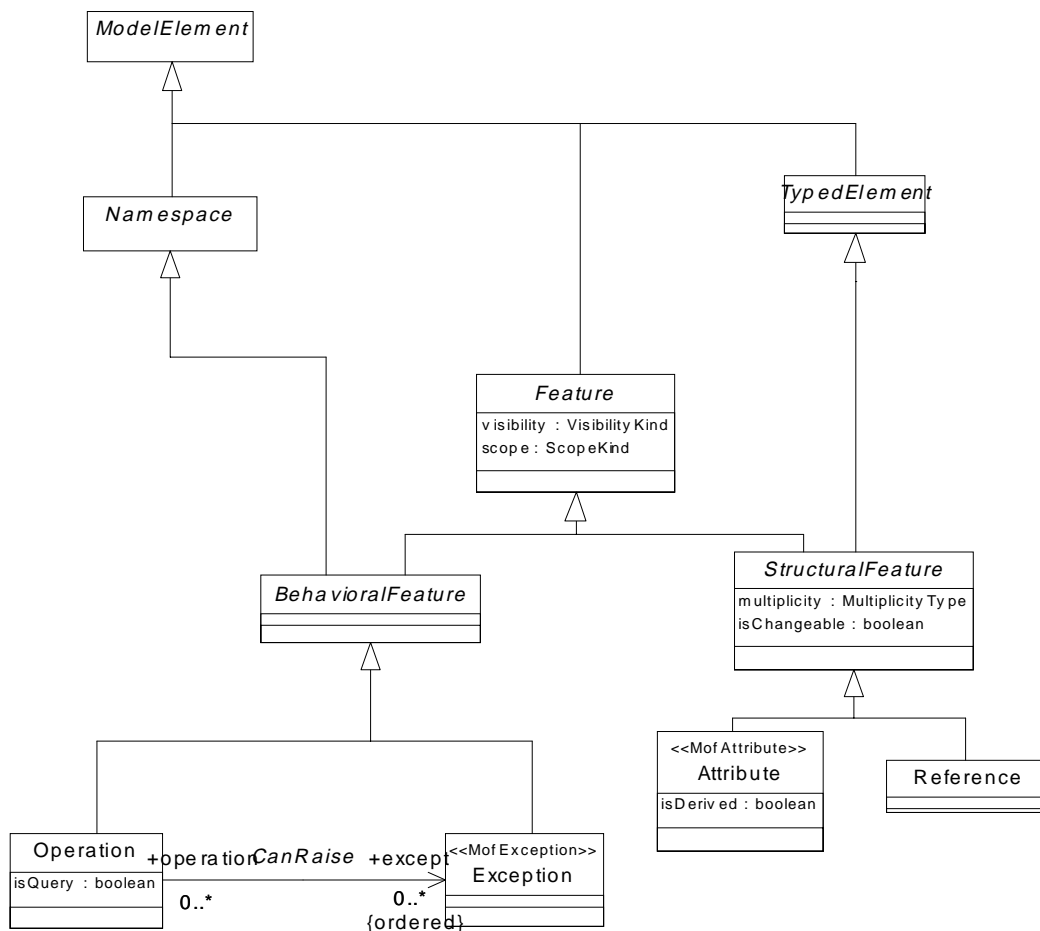


Figure 3-4 A TypeAlias relates a Classifier and its use in a DataType

As “Attribute” on page 3-40 describes, the type of a non-aggregate Class-valued Attribute is represented by a DataType with an “objRef” TypeCode. This is shown on the right hand side of Figure 3-4. Notice that D1’s “typeCode” value is an “objRef”

TypeCode with two components. The first component is a simple name for the Class, and the second is the repositoryId value for the TypeCode. While a repositoryId with the “local” syntax has been used here, any legal repositoryId syntax may be used.

The linkage from the use of a Classifier within the typeCode of a DataType has to be made by a round-about route. Any DataType instance may “contain” one or more TypeAlias instances, each of which has a “name” and a “type.” When the “name” of a TypeAlias is equal to a repositoryId in the DataType’s “typeCode,” the Classifier that is the TypeAlias’ “type” is the one that defines the TypeCode that the repositoryId belongs to.

If a TypeAlias related repositoryId belongs to a TypeCode that is embedded in the DataType’s “typeCode,” this means that the Classifier it relates to defines a component of the DataType (e.g., a member type of a struct or the content type of an array). A complex DataType may contain multiple TypeAliases.

SuperClasses

TypedElement

IDL

```
interface TypeAliasClass : TypedElementClass {
  readonly attribute TypeAliasUList all_of_type_type_alias;
  readonly attribute TypeAliasUList all_of_class_type_alias;

  TypeAlias create_type_alias (
    /* from ModelElement */           in ::Model::NameType name,
    /* from ModelElement */           in ::Model::AnnotationType annotation)
  raises (Reflective::MofError);
}; // end of interface TypeAliasClass

interface TypeAlias : TypeAliasClass, TypedElement {};
}; // end of interface TypeAlias
```

3.4.9 Feature Class

Feature

abstract

A Feature defines a characteristic of the ModelElement that contains it. Specifically, Classifiers are defined largely by a composition of Features. The Feature Class and its sub-Classes are illustrated in Figure 3-5.

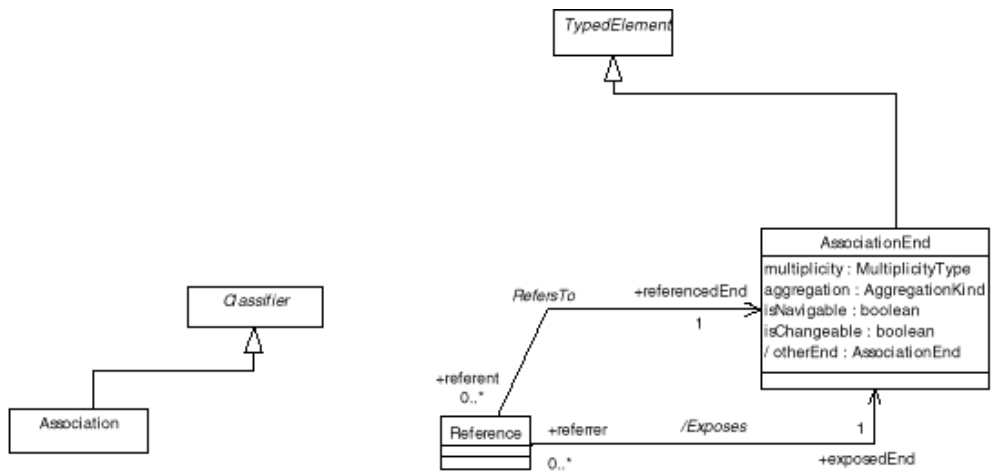


Figure 3-5 Feature Classes of the MOF Model

SuperClasses

ModelElement

Attributes

visibility

In the future, this Attribute will be used to limit the ability of ModelElements outside of this Feature's container to make use of it; see Section 3.6.6, "VisibilityKind," on page 3-95. The rules of visibility of MOF ModelElements are not currently specified.

<i>type:</i>	VisibilityKind
<i>multiplicity:</i>	exactly one

scope

The scope defines whether a Feature supports the definition of instances of the Classifier owning the Feature or of the Classifier as a whole. When scope is `instanceLevel`, the Feature is accessed through instances of the Feature's owning Classifier; when scope is `classifier`, the Feature is accessed through the Classifier itself (or through its instances). For StructuralFeatures, a scope of `instanceLevel` indicates that a value represented by the StructuralFeature is associated with each instance of the Classifier; a scope of `classifierLevel` indicates that the StructuralFeature value is shared by the Classifier and all its instances.

<i>type:</i>	ScopeKind
<i>multiplicity:</i>	exactly one

IDL

```
interface FeatureClass : ModelElementClass {
    readonly attribute FeatureUList all_of_type_feature;
}; // end of interface FeatureClass
```

```
interface Feature : FeatureClass, ModelElement {
    ScopeKind scope ()
        raises (Reflective::MofError);
    void set_scope (in ScopeKind new_value)
        raises (Reflective::MofError);
    VisibilityKind visibility ()
        raises (Reflective::MofError);
    void set_visibility (in VisibilityKind new_value)
        raises (Reflective::MofError);
};
```

*3.4.10 StructuralFeature**StructuralFeature**abstract*

A StructuralFeature defines a static characteristic of the ModelElement that contains it. The attributes and references of a Class define structural properties, which provide for the representation of the state of its instances.

SuperClasses

Feature, TypedElement

Attributes

multiplicity

Multiplicity defines constraints on the collection of instances or values that a StructuralFeature can hold. Multiplicity defines a lower and upper bound to the cardinality of the collection, although the upper bound can be specified as Unbounded. Additionally multiplicity defines two other characteristics of the collection: 1) a constraint on collection member ordering, and 2) a constraint on collection member uniqueness.

Specifically, Multiplicity contains an isOrdered field. When isOrdered is true, then the ordering of the elements in the set are preserved. Typically, a mechanism is provided for adding elements to the collection positionally. Multiplicity also has an isUnique field. When isUnique is true, then the collection is constrained to hold no more than one of any value or instance.

<i>type:</i>	MultiplicityType
<i>multiplicity:</i>	exactly one

isChangeable

The isChangeable attribute places restrictions on the use of certain operations, which could change the set of values or instances of the StructuralFeature, and on the operations that will get generated in IDL or other language generation. For any elaboration, no means are automatically created which provides a means of altering the attribute value. When IDL is generated, for instance, the operations, which are normally generated for changing the StructuralFeature, will not be generated. However, isChangeable does not actually constrain the StructuralFeature to make it immutable. Any operations explicitly defined in a model may change the StructuralFeature values or instances (assuming the operation would have otherwise been able to do so).

<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

IDL

```
interface StructuralFeatureClass : FeatureClass, TypedElementClass {
    readonly attribute StructuralFeatureUList all_of_type_structural_feature;
}; // end of interface StructuralFeatureClass
```

```
interface StructuralFeature : StructuralFeatureClass, Feature, TypedElement {
    MultiplicityType multiplicity ()
    raises (Reflective::MofError);
    void set multiplicity (in MultiplicityType new_value)
    raises (Reflective::MofError),
    boolean is_changeable ()
```

```

    raises (Reflective::MofError);
void set_is_changeable (in boolean new_value)
    raises (Reflective::MofError);
};

```

3.4.11 Attribute

Attribute

idl_substitute_name “MofAttribute”

An Attribute (referred to as a MofAttribute in the mapped IDL) defines a StructuralFeature which contains values for Classifiers or their instances.

SuperClasses

StructuralFeature

Contained Elements

None (not a Namespace)

Attributes

isDerived

A derived attribute is one whose values are not part of the state of the object instance, but whose values can be determined or computed. In a sense, all attributes are derived, since it is up to the class's implementation to hold or calculate the values. However, by convention, *isDerived* indicates that the derived state is based on other information in the model. Modification of the derived Attribute causes the information upon which the Attribute is derived to be updated.

<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

IDL

```

interface MofAttributeClass : StructuralFeatureClass {
    readonly attribute MofAttributeUList all_of_type_mof_attribute;
    readonly attribute MofAttributeUList all_of_class_mof_attribute;

    MofAttribute create_mof_attribute (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Feature */                 in ::Model::ScopeKind scope,
        /* from Feature */                 in ::Model::VisibilityKind visibility,
        /* from StructuralFeature */       in ::Model::MultiplicityType multiplicity,

```



```

/* from StructuralFeature */      in boolean is_changeable,
/* from MofAttribute */          in boolean is_derived)
  raises (Reflective::MofError);
}; // end of interface MofAttributeClass

interface MofAttribute : MofAttributeClass, StructuralFeature {
  boolean is_derived ()
    raises (Reflective::MofError);
  void set_is_derived (in boolean new_value)
    raises (Reflective::MofError);
};

```

3.4.12 Reference

A Reference defines a Classifier's knowledge of, and access to, links and their instances defined by an Association. Although a Reference derives much of its state from a corresponding AssociationEnd, it provides additional information; therefore, the MOF cannot adequately represent some meta-models without this mechanism. The inherited attributes defined in StructuralFeature (multiplicity and is_changeable) are constrained to match the values of its corresponding AssociationEnd. However, it has its own visibility, name, and annotation defined. For further discussion on Reference, its purpose, and how it derives its attributes, see Section 3.2.2, "Associations," on page 3-8.

Note – When creating a Reference, values for the inherited attributes of multiplicity and is_changeable must be supplied. These must be the same as the corresponding attributes on the AssociationEnd to which the Reference will subsequently be linked.

SuperClasses

StructuralFeature

References

referencedEnd

The referencedEnd of a Reference is the end representing the set of LinkEnds of principle interest to the Reference. The Reference provides access to the instances of that AssociationEnd's class, which are participants in that AssociationEnd's Association, connected through that AssociationEnd's LinkEnds. In addition, the Reference derives the majority of its state information - multiplicity, etc., from that Reference.	
<i>class:</i>	AssociationEnd

<i>defined by:</i>	RefersTo::referencedEnd
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	yes

exposedEnd

The exposedEnd of a Reference is the AssociationEnd representing the end of the Reference's owning Classifier within the defining Association.	
<i>class</i>	AssociationEnd
<i>defined by:</i>	Exposes::exposedEnd
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	yes

Constraints

The multiplicity for a Reference must be the same as the multiplicity for the referenced AssociationEnd. [C-21]

Classifier scoped References are not meaningful in the current M1 level computational model. [C-22]

A Reference can be changeable only if the referenced AssociationEnd is also changeable. [C-23]

The type attribute of a Reference and its referenced AssociationEnd must be the same. [C-24]

A Reference is only allowed for a navigable AssociationEnd [C-25]

The containing Class for a Reference must be equal to or a subtype of the type of the Reference's exposed AssociationEnd. [C-26]

The referenced AssociationEnd for a Reference must be visible from the Reference. [C-27]

IDL

```
interface ReferenceClass : StructuralFeatureClass {
  readonly attribute ReferenceUList all_of_type_reference;
  readonly attribute ReferenceUList all_of_class_reference;
  const string REFERENCE_MULTIPLICITY_MUST_MATCH_END =
    "org.omg.mof:constraint.model.reference.reference_multiplicity_must_match_end";
  const string REFERENCE_MUST_BE_INSTANCE_SCOPED =
    "org.omg.mof:constraint.model.reference.reference_must_be_instance_scoped";
  const string CHANGEABLE_REFERENCE_MUST_HAVE_CHANGEABLE_END =
    "org.omg.mof:constraint.model.reference.changeable_reference_must_have_changeable_end";
  const string REFERENCE_TYPE_MUST_MATCH_END_TYPE =
    "org.omg.mof:constraint.model.reference.reference_type_must_match_end_type";
```

```

const string REFERENCED_END_MUST_BE_NAVIGABLE =
"org.omg.mof:constraint.model.reference.referenced_end_must_be_navigable";
const string CONTAINER_MUST_MATCH_EXPOSED_TYPE =
"org.omg.mof:constraint.model.reference.container_must_match_exposed_type";
const string REFERENCED_END_MUST_BE_VISIBLE =
"org.omg.mof:constraint.model.reference.referenced_end_must_be_visible";

Reference create_reference (
/* from ModelElement */           in ::Model::NameType name,
/* from ModelElement */           in ::Model::AnnotationType annotation,
/* from Feature */                 in ::Model::ScopeKind scope,
/* from Feature */                 in ::Model::VisibilityKind visibility,
/* from StructuralFeature */       in ::Model::MultiplicityType multiplicity,
/* from StructuralFeature */       in boolean is_changeable)
    raises (Reflective::MofError);
}; // end of interface ReferenceClass

interface Reference : ReferenceClass, StructuralFeature {
    AssociationEnd exposed_end ()
        raises (Reflective::MofError);
    void set_exposed_end (in AssociationEnd new_value)
        raises (Reflective::MofError);
    AssociationEnd referenced_end ()
        raises (Reflective::MofError);
    void set_referenced_end (in AssociationEnd new_value)
        raises (Reflective::MofError);
};

```

3.4.13 BehavioralFeature Class

BehavioralFeature

abstract

A BehavioralFeature defines a dynamic characteristic of the ModelElement that contains it. Because a BehavioralFeature is partially defined by the Parameters it contains, it is both a Feature and a Namespace.

SuperClasses

Feature, Namespace

IDL

```

interface BehavioralFeatureClass : FeatureClass, NamespaceClass {
    readonly attribute BehavioralFeatureUList
        all_of_type_behavioral_feature;
}; // end of interface BehavioralFeatureClass

```

```

interface BehavioralFeature :
    BehavioralFeatureClass, Feature , Namespace {};

```

3.4.14 Operation

An Operation defines a dynamic feature which offers a service. The behavior of an operation is activated through the invocation of the operation.

SuperClasses

BehavioralFeature

Contained Elements

Parameter, Constraint; see [C-28]

Attributes

isQuery

Defining an Operation with an isQuery value of true denotes that the behavior of the operation will not alter the state of the object. The state of a Classifier, for this definition, is the set of values of all of the Classifier's class-scope and instance-scope StructuralFeatures.

For instance, an Operation of a Class, defined with a scope of instance, will not change the values or instances of any instance-scope StructuralFeature of the Class instance, as a result of invoking this Operation. An Operation of a Class with a scope of classifier will not change the values or instances of any of the classifier-scope StructuralFeatures or instance-scope StructuralFeatures.

This attribute does not define a constraint enforced by the model, but rather a promise that the operation's implementation is expected to uphold. An operation which is not defined as isQuery equals false is not guaranteed to change the state of its object. The isQuery constraint does not proscribe any specific implementation, so long as the definition of isQuery above is observed.	
--	--

<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

References

exceptions

An Operation, upon encountering an error or other abnormal condition, may raise an Exception. The exceptions reference provides the Operation with the set of Exceptions it is allowed to raise.	
<i>class:</i>	Exception
<i>defined by:</i>	CanRaise::except
<i>multiplicity:</i>	zero or more, ordered

Constraints

An Operation may only contain Parameters, Constraints and Tags. [C-28]

An Operation may have at most one Parameter whose direction is “return”. [C-29]

The Exceptions raised by an Operation must be visible to the Operation. [C-30]

IDL

```

interface OperationClass : BehavioralFeatureClass {
    readonly attribute OperationUList all_of_type_operation;
    readonly attribute OperationUList all_of_class_operation;
    const string OPERATION_CONTAINMENT_RULES =
        "org.omg.mof.constraint.model.operation.operation_containment_rules";
    const string OPERATIONS_HAVE_AT_MOST_ONE_RETURN =
        "org.omg.mof.constraint.model.operation.operations_have_at_most_one_return";
    const string OPERATION_EXCEPTIONS_MUST_BE_VISIBLE =
        "org.omg.mof.constraint.model.operation.operation_exceptions_must_be_visible";

    Operation create_operation (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Feature */                 in ::Model::ScopeKind scope,
        /* from Feature */                 in ::Model::VisibilityKind visibility,
        /* from Operation */               in boolean is_query)
        raises (Reflective::MofError);
}; // end of interface OperationClass

interface Operation : OperationClass, BehavioralFeature {
    boolean is_query ()
        raises (Reflective::MofError);
    void set_is_query (in boolean new_value)
        raises (Reflective::MofError);
    MofExceptionUList exceptions ()
        raises (Reflective::MofError);
    void set_exceptions (in MofExceptionUList new_value)
        raises (Reflective::MofError);
    void add_exceptions (in MofException new_element)

```

```

    raises (Reflective::MofError);
void add_exceptions_before (in MofException new_element, in MofException before_element)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_exceptions (in MofException old_element,
    in MofException new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_exceptions (in MofException old_element)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.4.15 Exception

Exception

idl_substitute_name “MofException”

An Exception (referred to as a MofException in the mapped IDL) defines an error or other abnormal condition. The Parameters of an Exception hold a record of an occurrence of the exceptional condition.

SuperClasses

BehavioralFeature

Contained Elements

Parameter; see [C-31]

Constraints

An Exception may only contain Parameters and Tags. [C-31]

An Exception’s Parameters must all have the direction “out.” [C-32]

IDL

```

interface MofExceptionClass : BehavioralFeatureClass {
    readonly attribute MofExceptionUList all_of_type_mof_exception;
    readonly attribute MofExceptionUList all_of_class_mof_exception;
    const string EXCEPTION_CONTAINMENT_RULES =
        "org.omg.mof.constraint.model.mof_exception.exception_containment_rules";
    const string EXCEPTIONS_HAVE_ONLY_OUT_PARAMETERS =
        "org.omg.mof.constraint.model.mof_exception.exceptions_have_only_out_parameters";

    MofException create_mof_exception (
        /* from ModelElement */ in ::Model::NameType name,
        /* from ModelElement */ in ::Model::AnnotationType annotation,
        /* from Feature */      in ::Model::ScopeKind scope,
        /* from Feature */      in ::Model::VisibilityKind visibility)
        raises (Reflective::MofError);
}; // end of interface MofExceptionClass

```

```
interface MofException : MofExceptionClass, BehavioralFeature {};
```

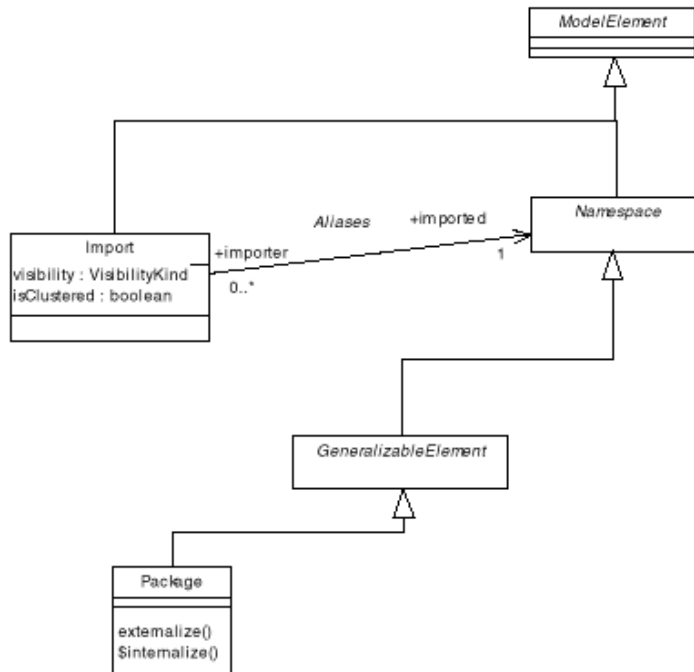


Figure 3-6 MOF Model Elements for Association

3.4.16 Association

An association defines a classification over a set of links, through a relationship between Classifiers. Each link which is an instance of the association denotes a connection between object instances of the Classifiers of the Association. The MOF restricts associations to binary, restricting each link to two participating objects. This restriction also means that the association is defined between two Classifiers (which may be the same Classifier). The name of the Association is considered directional if it provides a clearer or more accurate representation of the association when stated with one participating class first rather than the other. For instance, Operation CanRaise Exception is correct; Exception CanRaise Operation is incorrect.

The definition of an Association requires two AssociationEnds. If the name of the association is directional, the name is understood to read in the order: first contained element; association name; second contained element. These contained elements are AssociationEnd instances, and the reading of the subject; verb; object uses either the AssociationEnd name or the AssociationEnd's class name. The onus is on the MOF user to determine whether the name is directional, and to place the AssociationEnds in proper order within the Association's contents to support the name direction. The representation of a Classifier's knowledge of its participation in an association requires the use of a Reference.

SuperClasses

Classifier

Contained Elements

AssociationEnd, Constraint; see [C-33]

Attributes

isDerived

A derived association has no Links as instances. Instead, its Links are derived from other information in a meta-model. The addition, removal, or modification of a derived Association's Link causes the information upon which the Association is derived to be updated. The results of such an update are expected to appear, upon subsequent access of the derived Association's Links, to have the same effect as an equivalent operation on an Association which is not derived.	
<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

Constraints

An Association may only contain AssociationEnds, Constraints and Tags. [C-33]

Inheritance / generalization is not applicable to Associations. [C-34]

The values for “isLeaf” and “isRoot” on an Association must be true. [C-35]

An Association cannot be abstract. [C-36]

Associations must have visibility of “public.” [C-37]

An Association must be binary (i.e., it must have exactly two AssociationEnds). [C-38]

IDL

```

interface AssociationClass : ClassifierClass {
    readonly attribute AssociationUList all_of_type_association;
    readonly attribute AssociationUList all_of_class_association;
    const string ASSOCIATIONS_CONTAINMENT_RULES =
        "org.omg.mof.constraint.model.association.associations_containment_rules";
    const string ASSOCIATIONS_HAVE_NO_SUPERTYPES =
        "org.omg.mof.constraint.model.association.associations_have_no_supertypes";
    const string ASSOCIATIONS_MUST_BE_ROOT_AND_LEAF =
        "org.omg.mof.constraint.model.association.associations_must_be_root_and_leaf";
    const string ASSOCIATIONS_CANNOT_BE_ABSTRACT =
        "org.omg.mof.constraint.model.association.associations_cannot_be_abstract";
    const string ASSOCIATIONS_MUST_BE_PUBLIC =

```



```

"org.omg.mof:constraint.model.association.associations_must_be_public";
const string ASSOCIATIONS_MUST_BE_BINARY =
"org.omg.mof:constraint.model.association.associations_must_be_binary";

Association create_association (
    /* from ModelElement */           in ::Model::NameType name,
    /* from ModelElement */           in ::Model::AnnotationType annotation,
    /* from GeneralizableElement */   in boolean is_root,
    /* from GeneralizableElement */   in boolean is_leaf,
    /* from GeneralizableElement */   in boolean is_abstract,
    /* from GeneralizableElement */   in ::Model::VisibilityKind visibility,
    /* from Association */             in boolean is_derived)
    raises (Reflective::MofError);
}; // end of interface AssociationClass

interface Association : AssociationClass, Classifier {

    boolean is_derived ()
        raises (Reflective::MofError);
    void set_is_derived (in boolean new_value)
        raises (Reflective::MofError);
};

```

3.4.17 AssociationEnd

An association is composed of two AssociationEnds. Each AssociationEnd defines a Classifier participant in the Association, the role it plays, and constraints on sets of the Classifier instances participating. An instance of an AssociationEnd is a LinkEnd, which defines a relationship between a link, in instance of an Association, and an instance of the AssociationEnd's Classifier, provided in its type attribute.

SuperClasses

TypedElement

Attributes

multiplicity

Multiplicity defines constraints on sets of instances. Each instance of the Classifier defined by the opposite AssociationEnd's type defines a set which this multiplicity attribute constrains. Given one of those instances, x, the set is defined as the instances connected by LinkEnds of this AssociationEnd to that instance x. Refer to Section 3.6.5, "MultiplicityType," on page 3-78 for a description on how the multiplicity attribute constrains a set. In its use is describing AssociationEnds, isUnique has been constrained to be true, as a simplification. This constraint means that the same two instances cannot participate in more than one Link while participating under the same AssociationEnd. Normally, two instances cannot be linked by more than one Link of an Association at all. But when the AssociationEnd types allow the two instances switch ends, they can form a second Link without violating the isUnique constraint.

<i>type:</i>	MultiplicityType
<i>multiplicity:</i>	exactly one

aggregation

Certain associations define aggregations - directed associations with additional semantics (see Section 4.8, "Aggregation Semantics," on page 4-17). When an AssociationEnd is defined as composite or shared, the instance at "this" end of a Link is the composite or aggregate, and the instance at the "other" end is the component or subordinate..

<i>type:</i>	AggregationKind
<i>multiplicity:</i>	exactly one

isNavigable

The isNavigable attribute determines whether or not the AssociationEnd supports link "navigation". This has two implications:

- A Class defined with an appropriate Reference supports navigation of links from one Class instance to another. If isNavigable is false for an AssociationEnd, no such References may be created.
- Setting isNavigable to false also suppress a mapping's mechanisms for indexing links based on this AssociationEnd.

<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

isChangeable

The *isChangeable* attribute restricts the capability to perform actions that would modify sets of instances corresponding to this *AssociationEnd* (the same sets to which multiplicity is applied). Specifically, the set may be created when the instance defining the set - the instance at the opposite end of the Links - is created. This attribute does not make the set immutable. Instead, it affects the generation of operations in Model Elaboration which would allow modification of the set. For IDL generation, the only operation that allows the set to be modified would be one or more factory operations that create the instance and create the set. The modeler is free to define specific operations that allow modification of the set. Note that defining this *AssociationEnd* with *isChangeable* equals *false* places restrictions on the changeability of the other *AssociationEnd*, due to their interdependence..

<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

Operations

otherEnd

Provides the other <i>AssociationEnd</i> (i.e., not this one) in the enclosing <i>Association</i> .	
<i>return type:</i>	<i>AssociationEnd</i>
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>operation semantics:</i>	[S-2]

Constraints

The type of an *AssociationEnd* must be *Class*. [C-39]

The “isUnique” flag in an *AssociationEnd*’s multiplicity must be true. [C-40]

An *Association* cannot have two *AssociationEnds* marked as “ordered.” [C-41]

An *Association* cannot have an aggregation semantic specified for both *AssociationEnds*. [C-42]

IDL

```
interface AssociationEndClass : TypedElementClass {
  readonly attribute AssociationEndUList all_of_type_association_end;
  readonly attribute AssociationEndUList all_of_class_association_end;
  const string END_TYPE_MUST_BE_CLASS =
    "org.omg.mof.constraint.model.association_end.end_type_must_be_class";
  const string ENDS_MUST_BE_UNIQUE =
    "org.omg.mof.constraint.model.association_end.ends_must_be_unique";
```

```
const string CANNOT_HAVE_TWO_ORDERED_ENDS =
  "org.omg.mof:constraint.model.association_end.cannot_have_two_ordered_ends";
const string CANNOT_HAVE_TWO_AGGREGATE_ENDS =
  "org.omg.mof:constraint.model.association_end.cannot_have_two_aggregate_ends";

AssociationEnd create_association_end (
  /* from ModelElement */           in ::Model::NameType name,
  /* from ModelElement */           in ::Model::AnnotationType annotation,
  /* from AssociationEnd */         in boolean is_navigable,
  /* from AssociationEnd */         in ::Model::AggregationKind aggregation,
  /* from AssociationEnd */         in ::Model::MultiplicityType multiplicity,
  /* from AssociationEnd */         in boolean is_changeable)
raises (Reflective::MofError);
}; // end of interface AssociationEndClass

interface AssociationEnd : AssociationEndClass, TypedElement {
  boolean is_navigable ()
    raises (Reflective::MofError);
  void set_is_navigable (in boolean new_value)
    raises (Reflective::MofError);
  AggregationKind aggregation ()
    raises (Reflective::MofError);
  void set_aggregation (in AggregationKind new_value)
    raises (Reflective::MofError);
  MultiplicityType multiplicity ()
    raises (Reflective::MofError);
  void set_multiplicity (in MultiplicityType new_value)
    raises (Reflective::MofError);
  boolean is_changeable ()
    raises (Reflective::MofError);
  void set is_changeable (in boolean new_value);
  AssociationEnd other_end ()
    raises (Reflective::MofError);
}; // end of interface AssociationEnd
```

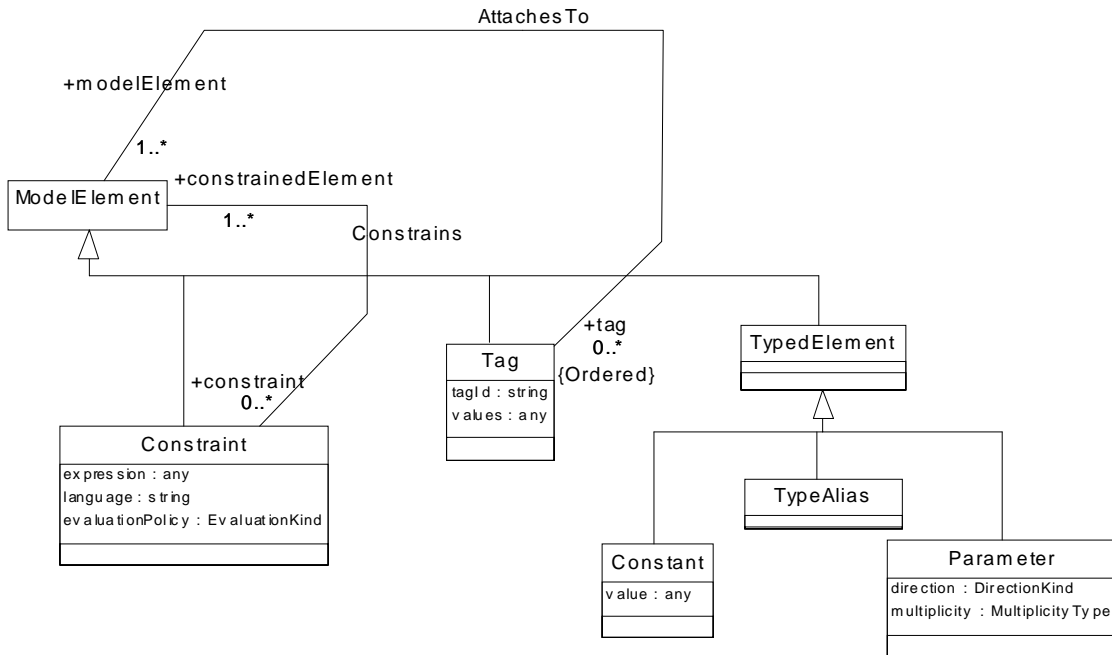


Figure 3-7 MOF Model Packaging

3.4.18 Package

A package is formed as a composition of `ModelElements`. A package defines a modeling unit, models are constructed and presented as packages. A model is a package. Packages are also used as organizational constructs in modeling. Nesting, importation, and generalization are used to manage the complexity of models.

SuperClasses

GeneralizableElement

Contained Elements

Package, Class, Association, DataType, Exception, Import, Constraint, Constant - see [C-43]

Operations

externalize

Externalize the Package and all of its ModelElements (transitive closure on the containment hierarchy) in a format specified by the format parameter, into a stream of type any.	
<i>return type:</i>	any
<i>isQuery:</i>	yes
<i>parameters:</i>	format : in FormatType
<i>exceptions:</i>	ObjectNotExternalizable, FormatNotSupported

internalize

Reify a model encoded in “stream” in some external format specified by “format” as a MOF Package.	
<i>return type:</i>	Package
<i>scope:</i>	classifier-scope
<i>parameters:</i>	format : in FormatType stream : in any
<i>exceptions:</i>	FormatNotSupported, IllformedExternalizedObject

Constraints

A Package may only contain Packages, Classes, DataTypes, Associations, Exceptions, Constraints, Imports and Tags. [C-43]

Packages cannot be declared as abstract. [C-44]

IDL

```

interface PackageClass : GeneralizableElementClass {
  readonly attribute PackageUList all_of_type_package;
  readonly attribute PackageUList all_of_class_package;
  const string PACKAGE_CONTAINMENT_RULES =
    "org.omg.mof.constraint.model.package.package_containment_rules";
  const string PACKAGES_CANNOT_BE_ABSTRACT =
    "org.omg.mof.constraint.model.package.packages_cannot_be_abstract";

  typedef string FormatType;

  exception FormatNotSupported {};
  exception ObjectNotExternalizable { string explanation; };
  exception IllformedExternalizedObject { string explanation; };

```

```

GeneralizableElement internalize (in PackageClass::FormatType format,
                                in any stream)
raises (PackageClass::FormatNotSupported,
        PackageClass::IllformedExternalizedObject,
        Reflective::MofError);
    Package create_package (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from GeneralizableElement */   in boolean is_root,
        /* from GeneralizableElement */   in boolean is_leaf,
        /* from GeneralizableElement */   in boolean is_abstract,
        /* from GeneralizableElement */   in ::Model::VisibilityKind visibility)
raises (Reflective::MofError);
}; // end of interface PackageClass

interface Package : PackageClass, GeneralizableElement {

any externalize (in PackageClass::FormatType format)
raises (PackageClass::ObjectNotExternalizable,
        PackageClass::FormatNotSupported,
        Reflective::MofError);
};

```

3.4.19 Import

An Import allows a Package to make use of ModelElements defined in some other Namespace. An Import object is related to another Namespace via the Aliases association. When a Package contains an Import object, it imports the associated Namespace. This means that ModelElements defined within the imported Namespace are visible in the importing Package.

An Import allows the visibility of the imported Package's contained ModelElements to be further restricted. An Import object represents either Package importing or Package clustering, depending on the “isClustered” attribute.

SuperClasses

ModelElement

Attributes

visibility

In the future, this Attribute will modify the visibility of imported ModelElements in the context of the importing Namespace. For a description of visibility kinds, see Section 3.6.6, “VisibilityKind,” on page 3-79. The MOF rules of visibility are not currently specified.	
--	--

<i>type:</i>	VisibilityKind
<i>multiplicity:</i>	exactly one

isClustered

The isClustered flags determines whether the Import object represents simple Package importation, or Package clustering.	
<i>type:</i>	boolean
<i>multiplicity:</i>	exactly one

*References**importedNamespace*

The Import knows about the Namespace that it references.	
<i>class:</i>	Namespace
<i>defined by:</i>	Aliases::imported
<i>multiplicity:</i>	exactly one

Constraints

The Namespace imported by an Import must be visible to the Import's containing Package. [C-45]

It is only legal for a Package to import or cluster Packages or Classes. [C-46]

Packages cannot import or cluster themselves. [C-47]

Packages cannot import or cluster Packages or Classes that they contain. [C-48]

Nested Packages cannot import or cluster other Packages or Classes. [C-49]

IDL

```

interface ImportClass : ModelElementClass {
    readonly attribute ImportUList all_of_type_import;
    readonly attribute ImportUList all_of_class_import;
    const string IMPORTED_NAMESPACE_MUST_BE_VISIBLE =
        "org.omg.mof:constraint.model.import.imported_namespace_must_be_visible";
    const string CAN_ONLY_IMPORT_PACKAGES_AND_CLASSES =
        "org.omg.mof:constraint.model.import.can_only_import_packages_and_classes";
    const string CANNOT_IMPORT_SELF =
        "org.omg.mof:constraint.model.import.cannot_import_self";
    const string CANNOT_IMPORT_NESTED_COMPONENTS =
        "org.omg.mof:constraint.model.import.cannot_import_nested_components";
    const string NESTED_PACKAGES_CANNOT_IMPORT =
        "org.omg.mof:constraint.model.import.nested_packages_cannot_import";

    Import create_import (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,

```



```

    /* from Import */
    /* from Import */
    raises (Reflective::MofError);
}; // end of interface ImportClass

```

```

interface Import : ImportClass, ModelElement {
    VisibilityKind visibility ()
    raises (Reflective::MofError);
    void set_visibility (in VisibilityKind new_value)
    raises (Reflective::MofError);
    boolean is_clustered ()
    raises (Reflective::MofError);
    void set_is_clustered (in boolean new_value)
    raises (Reflective::MofError);
    Namespace imported_namespace ()
    raises (Reflective::MofError);
    void set_imported_namespace (in Namespace new_value)
    raises (Reflective::MofError);
};

```

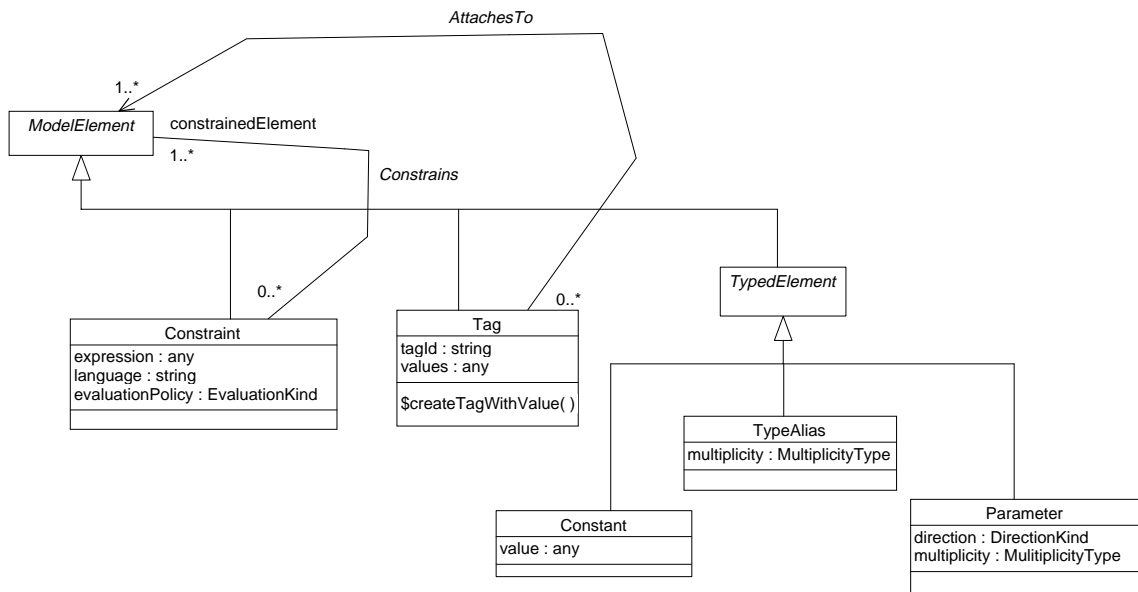


Figure 3-8 MOF Model - Other Elements

3.4.20 Parameter

A parameter provides a means of communication with operations and other BehavioralFeatures. A parameter passes or communicates values of its defined type.

SuperClasses

TypedElement

*Attributes**direction*

This attribute specifies the purpose of the parameter; to input a value, to output a value, both purposes, or to provide an operation return value.	
<i>type:</i>	DirectionKind
<i>multiplicity:</i>	exactly one

multiplicity

Multiplicity defines cardinality constraints on the set of instances or values that a Parameter can hold. Multiplicity defines a lower and upper bound on the set, although the upper bound can be specified as Unbounded. Additionally, multiplicity defines two other characteristics of the set: 1) constraints on set member ordering, and 2) constraints on unique set elements. Specifically, Multiplicity contains an isOrdered field. When isOrdered is true, the ordering of the elements in the set are preserved. Multiplicity also has an isUnique field. When isUnique is true, the set is constrained to hold no more than one of any value or instance.	
<i>type:</i>	MultiplicityType
<i>multiplicity:</i>	exactly one

IDL

```

interface ParameterClass : TypedElementClass {
    readonly attribute ParameterUList all_of_type_parameter;
    readonly attribute ParameterUList all_of_class_parameter;

    Parameter create_parameter (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Parameter */               in ::Model::DirectionKind direction,
        /* from Parameter */               in ::Model::MultiplicityType multiplicity)
        raises (Reflective::MofError);
}; // end of interface ParameterClass

interface Parameter : ParameterClass, TypedElement {
    DirectionKind direction ()
        raises (Reflective::MofError);
    void set_direction (in DirectionKind new_value)
        raises (Reflective::MofError);
    MultiplicityType multiplicity ()

```

```

    raises (Reflective::MofError);
void set_multiplicity (in MultiplicityType new_value)
    raises (Reflective::MofError);
};

```

3.4.21 Constraint

A Constraint defines a rule that restricts the state or behavior of one or more elements in the meta-model. When a Constraint is attached to a ModelElement, the rule it encodes applies to all relevant instances of the ModelElement in a model.

A Constraint rule, represented by the “expression” attribute, may be encoded in any form. The “language” attribute may be used to denote the language and encoding scheme used.

While some Constraints on a model may need to be treated as invariant, it is often convenient for other Constraints to be relaxed, for instance while a model is being edited. While, the “evaluationPolicy” attribute is used to represent these two cases, this information is at best advisory, since the MOF specification does not currently state how and when Constraints should be enforced.

Note – A Constraint cannot over-ride structural integrity rules defined by other parts of a meta-model (e.g., multiplicity specifications) or the integrity rules defined by a particular mapping of the meta-model to implementation technology.

SuperClasses

ModelElement

Attributes

expression

<p>The Constraint's expression attribute provides a representation of the constraint. The MOF has no specific requirement to interpret this expression, or to validate it against the language attribute. The specific handling of the expression will necessarily vary with the language used. However, it is expected that for any language an implementation accepts, it will enforce the constraints expressed in that language. The expression can be represented in any format, including text or a composition of objects.</p>	
<i>type</i>	any
<i>multiplicity:</i>	exactly one

language

The language representing this Constraint's expression is defined in this attribute. Since it is a string, most any language can be represented, including format variances in a language (e.g., OCL as text verses OCL as a parse tree).	
<i>type:</i>	string
<i>multiplicity:</i>	exactly one

evaluationPolicy

<p>Each constraint can be defined as immediate or deferred. For immediate Constraints, the constraint violation will be detected and reported within an operation in the chain of operations between the operation initiated by the MOF user and the operation that caused the constraint violation. The effect of an operation that violates an immediate constraint on the state of the object or objects being altered is implementation specific, and possibly undefined. However, if possible, an implementation should reverse the effects of the operation.</p> <p>For deferred Constraints, the constraint violation can only be detected when the Constraint is explicitly evaluated. The MOF defines an operation for such constraint evaluation, the verify operation. When the verify operation is invoked on a Constraint's container, the constraint is evaluated and a constraint violation is detected, if present.</p>	
<i>type:</i>	EvaluationKind
<i>multiplicity:</i>	exactly one

References**constrainedElements**

The Constraint has access to the ModelElements it constrains, through this reference. Note that the Constraint may depend on other ModelElements not represented in this reference. For instance, a Constraint may state that attribute A::x cannot exceed A::y in magnitude. The Constraint is on A::x, although it also depends on A::y. The link between the Constraint and A::x will be represented in the meta-model, not the link between the Constraint and A::y.	
<i>class:</i>	ModelElement
<i>defined by:</i>	Constrains::constrainedElement
<i>multiplicity:</i>	one or more

Constraints

Constraints, Tags, Imports, TypeAliases, and Constants cannot be constrained. [C-50]

A Constraint can only constrain ModelElements that are defined by or inherited by its immediate container. [C-51]

IDL

```

interface ConstraintClass : ModelElementClass {
    readonly attribute ConstraintUList all_of_type_constraint;
    readonly attribute ConstraintUList all_of_class_constraint;
    const string CANNOT_CONSTRAIN_THIS_ELEMENT =
        "org.omg.mof.constraint.model.constraint.cannot_constrain_this_element";
    const string CONSTRAINTS_LIMITED_TO_CONTAINER =
        "org.omg.mof.constraint.model.constraint.constraints_limited_to_container";

    enum EvaluationKind {immediate, deferred};

    Constraint create_constraint (
        /* from ModelElement */ in ::Model::NameType name,
        /* from ModelElement */in ::Model::AnnotationType annotation,
        /* from Constraint */   in any expression,
        /* from Constraint */   in string language,
        /* from Constraint */   in ::Model::ConstraintClass::EvaluationKind
                                evaluation_policy)
        raises (Reflective::MofError);
}; // end of interface ConstraintClass

interface Constraint : ConstraintClass, ModelElement {
    any expression ()
        raises (Reflective::MofError);
    void set_expression (in any new_value)
        raises (Reflective::MofError);
    string language ()
        raises (Reflective::MofError);
    void set_language (in string new_value)
        raises (Reflective::MofError);
    ConstraintClass::EvaluationKind evaluation_policy ()
        raises (Reflective::MofError);
    void set_evaluation_policy (in ConstraintClass::EvaluationKind new_value)
        raises (Reflective::MofError);
    ModelElementSet constrained_elements ()
        raises (Reflective::MofError);
    void set_constrained_elements (in ModelElementSet new_value)
        raises (Reflective::MofError);
    void add_constrained_elements (in ModelElement new_element)
        raises (Reflective::MofError);
    void modify_constrained_elements (in ModelElement old_element,
                                      in ModelElement new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_constrained_elements (in ModelElement old_element)
        raises (Reflective::NotFound, Reflective::MofError);
};

```

3.4.22 Constant Class

This Class provides a mechanism for defining constant values of simple data types, in the support of model development.

SuperClasses

TypedElement

Attributes

value

This Attribute gives the value of the constant.	
<i>type</i>	LiteralType
<i>multiplicity:</i>	exactly one

Constraints

The type of a Constant and the type of its value must be the same. [C-52]

The type of a Constant must be a CORBA data type that is legal for a CORBA 2.3 constant declaration. [C-53]

IDL

```

interface ConstantClass : TypedElementClass {
  readonly attribute ConstantUList all_of_type_constant;
  readonly attribute ConstantUList all_of_class_constant;
  const string CONSTANTS_VALUE_MUST_MATCH_TYPE =
    "org.omg.mof.constraint.model.constant.constants_value_must_match_type";
  const string CONSTANTS_TYPE_MUST_BE_SIMPLE_DATA_TYPE =
    "org.omg.mof.constraint.model.constant.constants_type_must_be_simple_data_type";

  Constant create_constant (
    /* from ModelElement */           in ::Model::NameType name,
    /* from ModelElement */           in ::Model::AnnotationType annotation,
    /* from Constant */               in ::Model::LiteralType value)
    raises (Reflective::MofError);
}; // end of interface ConstantClass

interface Constant : ConstantClass, TypedElement {
  LiteralType value ()
    raises (Reflective::MofError);
  void set_value (in LiteralType new_value)
    raises (Reflective::MofError);
};

```

3.4.23 Tag

Tags provide a light-weight extension mechanism that allows mapping, vendor, and even customer specific information to be added to or associated with a meta-model. In essence, Tags are arbitrary name / value that can be attached to instances of most ModelElements.

A Tag has an attribute called “tagId” that denotes a category of meaning, and another attribute called “values” that parameterizes that meaning. Each Tag is related to one or more ModelElements by the AttachesTo Association. The Tag need not be contained within the meta-model of the ModelElement it “tags.”

The MOF specification does not generally define the values for the “tagId” or the application specific categories of meaning that they denote. The exception to this is Section 5.6, “Standard Tags for the IDL Mapping,” on page 5-35, which defines some Tags that can be used in a meta-model to “tune” the IDL produced for a meta-model by the mapping.

Since “tagId” values are not standardized, there is a risk that different vendors or user organizations will use the same values to denote different categories of meaning. If a “tagId” value is used to mean different things, problems can arise when meta-models using the value are exchanged.

To avoid such Tag collisions, it is recommended that “tagId” values should use the following scheme based on Java package naming. Each value should start with a prefix formed by reversing the Internet domain name of a “tagId” naming authority. This should be followed by a locally unique component. For instance, this might be a standard or product name followed by a name or names that denotes the meaning. Here are some examples:

```
"org.omg.mof.idl_prefix"  
"org.omg.mof.some_tag"  
"com.rational.rose.screen_position"  
"au.edu.dstc.elvin.event_type"
```

It is also recommended that “tagId” values should be spelled in all lower case using the underscore (“_”) character as a word separator.

Note – In defining new Tag categories, the meta-modeler should take account of the fact that the MOF Model has no Reference for navigating from a ModelElement to its attached Tags. This allows one to attach Tags to elements of a “frozen” meta-model. On the other hand, makes it harder for a “client” of the meta-model objects to find the Tags for an element. One option is to require relevant Tags to be Contained by the elements they AttachTo, or their parents.

SuperClasses

ModelElement

*Attributes**tagId*

Gives the category of meaning for the Tag. The values for this attribute and their associated meanings are not standardized here. See discussion above.	
<i>type:</i>	string
<i>multiplicity:</i>	exactly one

values

Carries additional information (e.g. “parameters”) associated with the Tag.	
<i>type:</i>	any
<i>multiplicity:</i>	zero or more; not ordered; not unique

*References**elements*

The ModelElement or ModelElements that this Tag is attached to.	
<i>class:</i>	ModelElement
<i>defined by:</i>	AttachesTo::modelElement
<i>multiplicity:</i>	one or more

IDL

```

interface TagClass : ModelElementClass {
  readonly attribute TagUList all_of_type_tag;
  readonly attribute TagUList all_of_class_tag;

  Tag create_tag (
    /* from ModelElement */in ::Model::NameType name,
    /* from ModelElement */in ::Model::AnnotationType annotation,
    /* from Tag */          in string tag_id,
    /* from Tag */          in AnyBag values)
    raises (Reflective::MofError);
}; // end of interface TagClass

interface Tag : TagClass, ModelElement {
  string tag_id ()
    raises (Reflective::MofError);
  void set_tag_id (in string new_value)
    raises (Reflective::MofError);
  AnyBag values ()
    raises (Reflective::MofError);
}

```



```

void set_values (in AnyBag new_value)
  raises (Reflective::MofError);
void add_values (in any new_element)
  raises (Reflective::MofError);
void modify_values (in any old_element, in any new_element)
  raises (Reflective::NotFound, Reflective::MofError);
void remove_values (in any old_element)
  raises (Reflective::NotFound, Reflective::MofError);
ModelElementUList elements ()
  raises (Reflective::MofError);
void set_elements (in ModelElementUList new_value)
  raises (Reflective::MofError);
void add_elements (in ModelElement new_element)
  raises (Reflective::MofError);
void add_elements_before (in ModelElement new_element, in ModelElement before_element)
  raises (Reflective::NotFound, Reflective::MofError);
void modify_elements (in ModelElement old_element, in ModelElement new_element)
  raises (Reflective::NotFound, Reflective::MofError);
void remove_elements (in ModelElement old_element)
  raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5 MOF Model Associations

3.5.1 Contains

A meta-model is defined through a composition of ModelElements. A Namespace defines a ModelElement which composes other ModelElements. Since Namespace has several subclasses, there is a sizable combinatorial set of potential Namespace-ModelElement pairings. However, some of these pairings are not appropriate for building an object-oriented meta-model, such as a Class containing a Package (see Section 3.3.4, “The MOF Model Containment Hierarchy,” on page 3-14. This approach factors the container mechanisms into one abstraction, and allows the greatest flexibility for future changes to the MOF Model.

Ends

container

Each Namespace is a composition of zero or more ModelElements.	
<i>class:</i>	Namespace
<i>multiplicity:</i>	zero or one
<i>aggregation:</i>	Namespace forms a composite aggregation of ModelElements

containedElement

Each ModelElement, with the exception of top-level packages participates in the association as a containedElement.	
<i>class:</i>	ModelElement
<i>multiplicity:</i>	zero or more; ordered

IDL

```

interface Contains : Reflective::RefAssociation {
  ContainsLinkSet all_contains_links ();
  boolean exists (in Namespace container, in ModelElement contained_element)
    raises (Reflective::MofError);
  Namespace container (in ModelElement contained_element)
    raises (Reflective::MofError);
  ModelElementUList contained_element (in Namespace container)
    raises (Reflective::MofError);
  void add (in Namespace container, in ModelElement contained_element)
    raises (Reflective::MofError);
  void add_before_contained_element (in Namespace container,
    in ModelElement contained_element,
    in ModelElement before)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_container (in Namespace container,
    in ModelElement contained_element,
    in Namespace new_container)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_contained_element (in Namespace container,
    in ModelElement contained_element,
    in ModelElement new_contained_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (in Namespace container, in ModelElement contained_element)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.2 Generalizes

The Association defined on GeneralizableElement. A Link of this Association represents a supertype/subtype relationship (or a generalizes/specializes relationship).

*Ends**supertype*

The GeneralizableElement that is more general is the supertype.	
<i>class:</i>	GeneralizabelElement
<i>multiplicity:</i>	zero or more (a GeneralizableElement may have zero or more supertypes); ordered

subtype

The subtype is the GeneralizableElement that is more specific. The supertype Generalizes the subtype.	
<i>class:</i>	GeneralizableElement
<i>multiplicity:</i>	zero or more (a GeneralizableElement may have zero or more subtypes)

IDL

```

interface Generalizes : Reflective::RefAssociation {
  GeneralizesLinkSet all_generalizes_links ();
  boolean exists (in GeneralizableElement supertype,
                 in GeneralizableElement subtype)
    raises (Reflective::MofError);
  GeneralizableElementUList supertype (in GeneralizableElement subtype)
    raises (Reflective::MofError);
  GeneralizableElementSet subtype (in GeneralizableElement supertype)
    raises (Reflective::MofError);
  void add (in GeneralizableElement supertype,
           in GeneralizableElement subtype)
    raises (Reflective::MofError);
  void add_before_supertype (in GeneralizableElement supertype,
                            in GeneralizableElement subtype,
                            in GeneralizableElement before)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_supertype (in GeneralizableElement supertype,
                        in GeneralizableElement subtype,
                        in GeneralizableElement new_supertype)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_subtype (in GeneralizableElement supertype,
                      in GeneralizableElement subtype,
                      in GeneralizableElement new_subtype)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (in GeneralizableElement supertype,
              in GeneralizableElement subtype)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.3 RefersTo

A Reference derives most of its state from the AssociationEnd that it is linked to, based on this Association. For a Class defined with a Reference, each of its instances can be used to access the referenced object or objects. Those referenced objects will be of the Class defined by this referencedEnd AssociationEnd, playing the defined end.

Ends

referent

The Reference which is providing the reference through which instances playing the end-defined by the AssociationEnd can be accessed.	
<i>class:</i>	Reference
<i>multiplicity:</i>	zero or more; not ordered (an AssociationEnd may or may not be used by any number of References).

referencedEnd

The AssociationEnd which provides the majority of information for the Reference, including the LinkEnds which supply the referenced instances.	
<i>class:</i>	AssociationEnd
<i>multiplicity:</i>	exactly one

IDL

```
interface RefersTo : Reflective::RefAssociation {
    RefersToLinkSet all_refers_to_links ();
    boolean exists (in Reference referent, in AssociationEnd referenced_end)
        raises (Reflective::MofError);
    ReferenceSet referent (in AssociationEnd referenced_end)
        raises (Reflective::MofError);
    AssociationEnd referenced_end (in Reference referent)
        raises (Reflective::MofError);
    void add (in Reference referent, in AssociationEnd referenced_end)
        raises (Reflective::MofError);
    void modify_referent (in Reference referent,
        in AssociationEnd referenced_end,
        in Reference new_referent)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_referenced_end (in Reference referent,
        in AssociationEnd referenced_end,
        in AssociationEnd new_referenced_end)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (in Reference referent, in AssociationEnd referenced_end)
        raises (Reflective::NotFound, Reflective::MofError);
};
```

3.5.4 Exposes

Exposes

derived

A Reference defines a reference for a Class. For an instance of that class, which holds one or more links to some object or objects conforming to the reference, the instance will be playing the role (end) defined by the AssociationEnd in this Association.

Although this association can be derived in the current MOF, the use of n-ary associations, where a single Class has multiple ends specification of this Association, is necessary.

Ends

referrer

The Reference that is providing the exposedEnd's class instances within the Reference's Classifier.	
<i>class:</i>	Reference
<i>multiplicity:</i>	zero or more; not ordered (an AssociationEnd may or may not be used by any number of References).
<i>changeable:</i>	yes

exposedEnd

The AssociationEnd representing the Reference's owning Classifier's end in the Association.	
<i>class:</i>	AssociationEnd
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	yes

Derivation

See [S-13]. For a given Reference, the Link of this Association is derived as follows:

- The referrer's Reference is the given Reference.
- The exposedEnd's AssociationEnd is the given Reference's referent's container Association's other AssociationEnd.

IDL

```
interface Exposes : Reflective::RefAssociation {
    ExposesLinkSet all_exposes_links ();
    boolean exists (in Reference referrer, in AssociationEnd exposed_end)
    raises (Reflective::MofError);
```

```

ReferenceSet referrer (in AssociationEnd exposed_end)
  raises (Reflective::MofError);
AssociationEnd exposed_end (in Reference referrer)
  raises (Reflective::MofError);
void add (in Reference referrer, in AssociationEnd exposed_end)
  raises (Reflective::MofError);
void modify_referrer (in Reference referrer,
  in AssociationEnd exposed_end,
  in Reference new_referrer)
  raises (Reflective::NotFound, Reflective::MofError);
void modify_exposed_end (in Reference referrer,
  in AssociationEnd exposed_end,
  in AssociationEnd new_exposed_end)
  raises (Reflective::NotFound, Reflective::MofError);
void remove (in Reference referrer, in AssociationEnd exposed_end)
  raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.5 *IsOfType*

A Link between a TypedElement subclass and a Classifier supports the definition of the TypedElement.

Ends

type

The type defining the TypedElement.	
<i>class:</i>	Classifier
<i>multiplicity:</i>	exactly one

typedElements

The set of typed elements supported by a Classifier.	
<i>class:</i>	TypedElement
<i>multiplicity:</i>	zero or more

IDL

```

interface IsOfType : Reflective::RefAssociation {
  IsOfTypeLinkSet all_is_of_type_links ();
  boolean exists (in Classifier type, in TypedElement typed_elements)
    raises (Reflective::MofError);
  Classifier type (in TypedElement typed_elements)
    raises (Reflective::MofError);
  TypedElementSet typed_elements (in Classifier type)
    raises (Reflective::MofError);
  void add (in Classifier type, in TypedElement typed_elements)

```

```

    raises (Reflective::MofError);
void modify_type (in Classifier type,
                 in TypedElement typed_elements,
                 in Classifier new_type)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_typed_elements (in Classifier type,
                           in TypedElement typed_elements,
                           in TypedElement new_typed_elements)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (in Classifier type, in TypedElement typed_elements)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.6 CanRaise

Relates Operations to the Exceptions that they can raise.

Ends

operation

Given an Exception, the set of Operations which can Raise that Exception.	
<i>class:</i>	Operation
<i>multiplicity:</i>	zero or more (an Exception may be defined that is not currently used by any Operation; an Exception may be raised by multiple Operations).

except

The set of Exceptions for an Operation.	
<i>class:</i>	Exception
<i>multiplicity:</i>	zero or more (an Operation may be defined to raise no exception, or multiple exceptions); ordered (an Operation's Exceptions are ordered).

IDL

```

interface CanRaise : Reflective::RefAssociation {
    CanRaiseLinkSet all_can_raise_links ();
    boolean exists (in ::Model::Operation operation, in MofException except)
        raises (Reflective::MofError);
    OperationSet operation (in MofException except)
        raises (Reflective::MofError);
    MofExceptionUList except (in ::Model::Operation operation)
        raises (Reflective::MofError);
    void add (in ::Model::Operation operation, in MofException except)
        raises (Reflective::MofError);
};

```

```

void add_before_except (in ::Model::Operation operation,
    in MofException except,
    in MofException before)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_operation (in Operation operation,
    in MofException except,
    in Operation new_operation)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_except (in ::Model::Operation operation,
    in MofException except,
    in MofException new_except)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (in ::Model::Operation operation, in MofException except)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.7 Aliases

An Import aliases or imports a single Namespace.

Ends

importer

A Namespace may be aliased by an Import, which is the importer.	
<i>class:</i>	Import
<i>multiplicity:</i>	zero or more (a Namespace may not be aliased, or may be aliased by multiple Imports).

imported

The Namespace that an Import imports or aliases.	
<i>class:</i>	Namespace
<i>multiplicity:</i>	exactly one

IDL

```

interface Aliases : Reflective::RefAssociation {
    AliasesLinkSet all_aliases_links ();
    boolean exists (in Import importer, in Namespace imported)
        raises (Reflective::MofError);
    ImportSet importer (in Namespace imported)
        raises (Reflective::MofError);
    Namespace imported (in Import importer)
        raises (Reflective::MofError);
    void add (in Import importer, in Namespace imported)
        raises (Reflective::MofError);
    void modify_importer (in Import importer,

```



```

        in Namespace imported,
        in Import new_importer)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_imported (in Import importer,
    in Namespace imported,
    in Namespace new_imported)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (in Import importer, in Namespace imported)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.8 Constrains

Each Constraint constrains one or more ModelElements.

Ends

constraint

A Constraint which constrains a ModelElement.	
<i>class:</i>	Constraint
<i>multiplicity:</i>	zero or more (a ModelElement need not be constrained, but could be constrained by more than one Constraint).

constrainedElement

The ModelElements that a Constraint holds its constraint against.	
<i>class:</i>	ModelElement
<i>multiplicity:</i>	one or more (a Constraint must constrain at least one ModelElement)

IDL

```

interface Constrains : Reflective::RefAssociation {
    ConstrainsLinkSet all_constrains_links ();
    boolean exists (in ::Model::Constraint constraint,
        in ModelElement constrained_element)
        raises (Reflective::MofError);
    ConstraintSet constraint (in ModelElement constrained_element);
    raises (Reflective::MofError)
    ModelElementSet constrained_element (in ::Model::Constraint constraint)
    raises (Reflective::MofError);
    void add (in ::Model::Constraint constraint,
        in ModelElement constrained_element)
        raises (Reflective::MofError);
    void modify_constraint (in ::Model::Constraint constraint,

```

```

        in ModelElement constrained_element,
        in Constraint new_constraint)
    raises (Reflective::NotFound, Reflective::MofError);
void modify_constrained_element (in ::Model::Constraint constraint,
    in ModelElement constrained_element,
    in ModelElement new_constrained_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove (in ::Model::Constraint constraint,
    in ModelElement constrained_element)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.5.9 DependsOn

DependsOn

derived

DependsOn is a derived Association that allows a client to identify the collection of ModelElements on which a given ModelElement *structurally* depends. The Association is derived from a number of other Associations in the MOF Model, as described below.

Note – The model of dependency that is embodied in in this Association is based solely on the structural relationships within a meta-model. In some cases, the structural dependencies have clear semantic parallels (e.g., the meaning of an Attribute depends on its type). In other cases the semantic parallel is more tenuous (e.g., a DataType only semantically depends on its container in the context of type identity).

Ends

dependent

This End is occupied by ModelElements that <i>structurally</i> depend on the ModelElement at the other End.	
<i>class:</i>	ModelElement
<i>multiplicity:</i>	zero or more (a ModelElement can have no ModelElement depend on it, or many may depend on it).
<i>changeable:</i>	no

provider

This End is occupied by ModelElements that have other ModelElements that <i>structurally</i> depend on them.	
<i>class:</i>	ModelElement
<i>multiplicity:</i>	zero or more (a ModelElement can depend on no other ModelElements or multiple ModelElements).
<i>changeable:</i>	no

Derivation

See [S-14]. A ModelElement (ME) depends on:

- “container” - its container Namespace from ModelElement::container
- “constraint” - any Constraints from ModelElement::constraints.
- “contents” - if ME is a Namespace, its contents from Namespace::contents.
- “specialization” - if ME is a GeneralizableElement, its supertypes from GeneralizableElement::supertypes.
- “import” if ME is an Import, the imported Package or Class from Import::importedNamespace
- “signature” - if ME is an Operation, the Exceptions it raises from Operation::exceptions.
- “type definition” - if ME is a TypedElement, the Classifier from TypedElement::type.
- “referenced ends” - if ME is a Reference, the two AssociationEnds from Reference::referencedEnd and Reference::exposedEnd
- “constrained elements” - if ME is a Constraint, the elements it constrains from Constraint::constrainedElements.
- “tagged elements” - if ME is a Tag, the elements it is attached to from Tag::elements.

IDL

```
interface DependsOn : Reflective::RefAssociation {
    DependsOnLinkSet all_depends_on_links ();
    boolean exists (in ModelElement dependent, in ModelElement provider)
        raises (Reflective::MofError);
    ModelElementSet dependent (in ModelElement provider)
        raises (Reflective::MofError);
    ModelElementSet provider (in ModelElement dependent)
        raises (Reflective::MofError);
};
```

3.5.10 AttachesTo

This association represents Tags attached to ModelElements. A ModelElement's Tags are ordered, although the ordering may not be of any significance, depending on the meaning of the Tags. Ordering is preserved in case some Tags, in conjunction with some defined semantics, requires an ordering.

Ends

modelElement

The ModelElements that an attached Tag describes, modifies, or otherwise associates.	
<i>class:</i>	ModelElement
<i>multiplicity:</i>	one or more (a Tag must be attached to at least one ModelElement).

tag

The set of Tags attached to a ModelElement.	
<i>class:</i>	Tag
<i>multiplicity:</i>	zero or more (a ModelElement need not have a Tag), ordered.

IDL

```
interface AttachesTo : Reflective::RefAssociation {
  AttachesToLinkSet all_attaches_to_links ();
  boolean exists (in ModelElement model_element, in ::Model::Tag tag)
    raises (Reflective::MofError);
  ModelElementSet model_element (in ::Model::Tag tag)
    raises (Reflective::MofError);
  TagUList tag (in ModelElement model_element)
    raises (Reflective::MofError);
  void add (in ModelElement model_element, in ::Model::Tag tag)
    raises (Reflective::MofError);
  void add_before_tag (in ModelElement model_element,
    in ::Model::Tag tag,
    in Tag before)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_model_element (in ModelElement model_element,
    in ::Model::Tag tag,
    in ModelElement new_model_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_tag (in ModelElement model_element,
    in ::Model::Tag tag,
    in Tag new_tag)
```

```

    raises (Reflective::NotFound, Reflective::MofError);
void remove (in ModelElement model_element, in ::Model::Tag tag)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

3.6 MOF Model Data Types

The following data types are part of the MOF Model. Each data type is represented in the MOF Model as an instance of the `DataType` class with a typecode that corresponds to the Data Type.

3.6.1 CORBA Basic Types

Each CORBA basic type is potentially available as an instance of `DataType`. The basic types used in the MOF Model are **boolean**, **long**, **unsigned long**, **any**, **string**, and **TypeCode**.

3.6.2 NameType

`NameType` is an unbounded string data type used to represent `ModelElement` names.

IDL

```
typedef string NameType;
```

3.6.3 AnnotationType

`AnnotationType` is an unbounded string data type used to represent `ModelElement` annotations.

IDL

```
typedef string AnnotationType;
```

3.6.4 TypeDescriptor

A `TypeDescriptor` value represents a particular type. By defining `TypeDescriptor` as a CORBA `TypeCode`, the MOF makes use of CORBA's types and type definition capabilities.

IDL

```
typedef CORBA::TypeCode TypeDescriptor;
```

3.6.5 MultiplicityType

MultiplicityType is a structure (record) type that is used to specify the multiplicity properties of an Attribute, Parameter, Reference, or AssociationEnd.

*Fields**lower*

This field gives the lower bounds on the number of elements allowed for the Attribute, Parameter, Reference, or AssociationEnd.	
---	--

<i>type:</i>	long
--------------	------

upper

This field gives the upper bounds on the number of elements allowed for the Attribute, Parameter, Reference or AssociationEnd. A value of Unbounded (see Section 3.8.1, “Unbounded,” on page 3-85) indicates that there is no upper bound on the number of elements.	
--	--

<i>type:</i>	long
--------------	------

isOrdered

This flag indicates whether the order of the elements corresponding to the Attribute, Parameter, Reference, or AssociationEnd has any semantic significance.	
--	--

<i>type:</i>	boolean
--------------	---------

isUnique

This flag indicates whether or not the elements corresponding to the Attribute, Parameter, Reference, or AssociationEnd are required (or guaranteed) to be unique.	
--	--

<i>type:</i>	boolean
--------------	---------

Constraints

The “lower” bound of an MultiplicityType to be “Unbounded.” [C-54]

The “lower” bound of a MultiplicityType cannot exceed the “upper” bound. [C-55]

The “upper” bound of a MultiplicityType cannot be less than 1. [C-56]

If a MultiplicityType specifies bounds of [0..1] or [1..1], the “is_ordered” and “is_unique” values must be false. [C-57]

IDL

```

struct MultiplicityType {
    long lower;
    long upper;
    boolean isOrdered;
    boolean isUnique;
};

const string LOWER_CANNOT_BE_NEGATIVE_OR_UNBOUNDED =
    "org.omg:constraint.model.multiplicity_type.lower_cannot_be_negative_or_unbounded";
const string LOWER_CANNOT_EXCEED_UPPER =
    "org.omg:constraint.model.multiplicity_type.lower_cannot_exceed_upper";
const string UPPER_MUST_BE_POSITIVE =
    "org.omg:constraint.model.multiplicity_type.upper_must_be_positive";
const string MUST_BE_UNORDERED_NONUNIQUE =
    "org.omg:constraint.model.multiplicity_type.must_be_unordered_nonunique";

```

3.6.6 VisibilityKind

This data type enumerates the three possible kinds of visibility for a ModelElement outside of its container. These are:

1. “public_vis,” which allows anything that can use ModelElement’s container to also use the ModelElement.
2. “protected_vis,” which allows use of the ModelElement within containers that inherit from this one’s container.
3. “private_vis,” which denies all outside access to the ModelElement.

Note – The rules governing visibility of ModelElements in the MOF are yet to be specified. As an interim measure, all ModelElements are deemed to be visible, irrespective of the “visibility” attribute settings. The IDL mapping specification includes minimal preconditions on visibility to ensure that generated IDL is compilable (see Section 5.5, “Preconditions for IDL Generation,” on page 5-33).

IDL

```

enum VisibilityKind {public_vis, private_vis, protected_vis};

```

3.6.7 DepthKind

DepthKind enumerates the two choices of depth semantic for the “verify” operation.

Container

ModelElement

*IDL***enum DepthKind {shallow, deep};***3.6.8 DirectionKind*

DirectionKind enumerates the possible directions of information transfer for Operation and Exception Parameters

*IDL***enum DirectionKind {in_dir, out_dir, inout_dir, return_dir};***3.6.9 ScopeKind*

ScopeKind enumerates the possible “scopes” for Attributes and Operations.

*IDL***enum ScopeKind {instance_level, classifier_level};***3.6.10 AggregationKind*

AggregationKind enumerates the possible aggregation semantics for Associations (specified via AssociationEnds).

Note – Aggregation semantics in the MOF is intended to be aligned with UML. Unfortunately, the OMG UML specification does not define the meaning of “shared” aggregation for UML. As an interim measure, the use of “shared” aggregation in MOF meta-models is discouraged.

*IDL***enum AggregationKind {none, shared, composite};***3.6.11 EvaluationKind*

EvaluationKind enumerates the possible models for Constraint evaluation.

Container

Constraint

IDL

```
enum EvaluationKind {immediate, deferred};
```

3.6.12 DependencyKind

DependencyKind is a string that represents the “cause” for a particular dependency as expressed by the DependsOn Association. The type declaration is accompanied by a small number of predefined dependency types.

When a ModelElement depends on a second model element under one kind of dependency; and the second model element depends on a third under some other kind of dependency; then the first ModelElement depends on the third ModelElement. However, the kind of dependency cannot be specified, based on the other two dependency kinds, except to categorize the dependency as indirect.

Container

ModelElement

IDL

```
typedef string DependencyKind;
```

3.6.13 FormatType

FormatType is a string type whose values denoted externalization formats.

Container

Package

IDL

```
typedef string FormatType;
```

3.6.14 LiteralType

LiteralType is the type used to represent the value of Constants.

IDL

typedef any LiteralType;

3.6.15 VerifyResultKind

VerifyResultKind enumerates the outcomes of the “verify” operation.

Container

ModelElement

IDL

enum VerifyResultKind {valid, invalid, published};

3.6.16 ViolationType

ViolationType is a structure (record) type that is used to return a description of an error detected during ModelElement verification. The fields have the same names and meanings as for the Reflective MofError exception. Refer to Section 5.4, “Exception Framework,” on page 5-24 for the complete description.

*Fields**errorKind*

This field will contain the kind string for the error being reported.	
<i>type:</i>	string

elementInError

This field will give the ModelElement for the instance which the error is being reported (e.g., the ModelElement to which a violated Constraint belongs, or AssociationEnd that is underflowed).	
<i>type:</i>	Reflective::RefObject

valuesInError

This field gives a value or values that caused the error.	
<i>type:</i>	Reflective::NamedValueList

errorDescription

This field gives a human intelligible textual description of the error	
<i>type:</i>	string

Container

ModelElement

IDL

```

struct Violation {
    string error_kind;
    Reflective::RefObject element_in_error;
    Reflective::NamedValueList values_in_error;
    string error_description;
};

```

3.7 MOF Model Exceptions

The following exceptions are contained in the MOF Model Package. The generated IDL interfaces for the MOF Model make use of more exceptions, which are defined in the Reflective Package (see the Reflective Type Packages chapter) and assigned to operations based on criteria determinable during generation.

3.7.1 NameNotFound

The NameNotFound exception is raised when a lookup of a simple name has failed.

parameters: name : out NameType

The name parameter gives the string value that could not be found in the Namespace or extended Namespace searched by the operation.

Container

Namespace

3.7.2 NameNotResolved

The NameNotResolved exception is raised when resolution of a qualified name has failed.

parameters: explanation : out string

 restOfName : out NameType (multiplicity: zero or more; ordered; not unique)

The `restOfName` parameter contains that part of the qualified name that was not resolved. The explanation parameter can have the following values with the corresponding interpretation:

- “InvalidName”: the first name in `restOfName` was malformed.
- “MissingName”: the first name in `restOfName` could not be resolved as no name binding exists for that name.
- “NotNameSpace”: the first name in `restOfName` did not resolve to a `NameSpace` when a `NameSpace` was expected.
- “CannotProceed”: the first name in `restOfName` could not be resolved (for any other reason).

Container

Namespace

3.7.3 ObjectNotExternalizable

An object cannot be externalized in the requested format.

parameters: explanation : out string

Container

Package

3.7.4 FormatNotSupported

The requested format for `internalize/externalize` is not supported.

parameters: none

Container

Package

3.7.5 IllformedExternalizedObject

The externalized form of the object does not conform to the format expected.

parameters: explanation : out string

Container

Package

3.8 MOF Model Constants

The following Constants form part of the MOF Model.

3.8.1 Unbounded

This constant is used in the context of MultiplicityType to represent an unlimited upper bound on a cardinality (see Section 3.6.5, “MultiplicityType,” on page 3-78).

Container

Model

IDL

```
const unsigned long UNBOUNDED = -1;
```

3.8.2 The Standard DependencyKinds

These constants (ContainerDep, ContentsDep, SignatureDep, ConstraintDep, ConstrainedElementsDep, SpecializationDep, ImportDep, TypeDefinitionDep, ReferencedEndsDep, TaggedElementsDep, IndirectDep, and AllDep) denote the standard dependency categories and pseudo-categories. Refer to “ModelElement” on page 3-15 and “DependsOn” on page 3-74 for detailed explanations.

Container

ModelElement

IDL

```
const DependencyKind CONTAINER_DEP = "container";
const DependencyKind CONTENTS_DEP = "contents";
const DependencyKind SIGNATURE_DEP = "signature";
const DependencyKind CONSTRAINT_DEP = "constraint";
const DependencyKind CONSTRAINED_ELEMENTS_DEP = "constrained elements";
const DependencyKind SPECIALIZATION_DEP = "specialization";
const DependencyKind IMPORT_DEP = "import";
const DependencyKind TYPE_DEFINITION_DEP = "type definition";
const DependencyKind REFERENCED_ENDS_DEP = "referenced ends";
const DependencyKind TAGGED_ELEMENTS_DEP = "tagged elements";
const DependencyKind INDIRECT_DEP = "indirect";
const DependencyKind ALL_DEP = "all";
```

3.9 MOF Model Constraints

3.9.1 MOF Model Constraints and other M2 Level Semantics

This section defines the semantic constraints that apply to the MOF Model. These are expressed as M2-level Constraints and are formally part of the MOF Model (i.e., they are a required part of a representation of the MOF Model as MOF meta-objects or in the MOF Model / XMI interchange format).

The section also provides OCL semantic specifications for most M2-level Operations, derived Attributes, and derived Associations in the MOF Model, and for a collection of “helper” functions used by them and the Constraints. These semantic specifications need not be present in a representation of the MOF Model. Indeed, this document does not specify how they should be represented.

Note – The use of OCL in the MOF Model specification does not imply a requirement to use OCL evaluation as part of a/n MOF Model server’s implementation. Furthermore, if that approach is used, it is anticipated that the implementor may rewrite the OCL rules to make evaluation more efficient. For example, the Constraint OCL could be rewritten as pre-conditions on the appropriate mapped update operations.

3.9.2 Notational Conventions

3.9.2.1 Notation for MOF Model Constraints

The M2-level Constraints on the MOF Model are described in the following notation:

[C-xxx]	ConstraintName
evaluation policy:	immediate or deferred
description:	brief english description

context SomeClassifierName

inv: . . .

The meaning of the above is as follows:

- “[C-xxx]” is the cross reference tag for the Constraint used elsewhere in this document.
- “ConstraintName” is the name for the Constraint in the MOF Model. The IDL mapping uses this name to produce the MofError “kind” string for the Constraint. These strings appear in the generated IDL for the MOF Model, as described in Section 5.8.17, “Constraint Template,” on page 5-91.
- The “evaluation policy” states whether the Constraint should be checked on any relevant update operation, or whether checking should be deferred until full meta-model validation is triggered. It defines the Constraint’s “evaluationPolicy” value.

- The “description” is a brief non-normative synopsis of the Constraint. It could be used as the Constraint’s “annotation” value.
- The OCL for the Constraint is defined using the OCL syntax defined in UML 1.3.

The OCL for the Constraints start with a “context” clause that names a ModelElement in the MOF Model. This serves two purposes:

1. It defines the context in which the OCL constraint should be evaluated (i.e., the M3-level Class or DataType whose instances are constrained by the OCL).
2. It defines the “constrainedElements” and “container” for the Constraint.

While the OCL for the Constraints are mostly expressed as invariants, this should not be taken literally. Instead, the Constraint OCL should be viewed as:

- a pre-condition on the relevant IDL operations for “immediate” Constraints, or
- a part of the specification of ModelElement’s “verify” Operation for “deferred” Constraints.

The Constraints in the MOF Model are expressed as restrictions on either Classes or DataTypes. Each one applies to (“Constrains”) a single Classifier, and each one is defined to be contained by the Classifier that it applies. The “language” attribute of each Constraint is either “MOF-OCL” (for those with complete OCL specifications) or “Other.” The “expression” attribute should be the normative OCL defined here, even if different (but equivalent) OCL is used in a/n MOF Model server’s implementation.

3.9.2.2 *Notation for Operations, derived Attributes and derived Association*

The semantics of M2-level Operations, derived Attributes, and derived Associations on the MOF Model are described in the following notation:

```
[O-xxx]      ModelElementName
kind:        classification
description:  brief english description

context ClassifierName::OperationName(...) : ...
post: result = ...

or

context ClassName::AttributeName() : ...
post: result = ...

or

context ClassName::ReferenceName() : ...
post: result = ...
```

The meaning of the above is as follows:

- “[O-xxx]” is the cross reference tag for the semantic description that may be used elsewhere in this document.
- “ModelElementName” is the name of the Attribute, Operation, or Association in the MOF Model whose semantics is described.

- The “classification” describes the kind of the ModelElement (e.g., “readonly derived Attribute” or “query Operation”).
- The “description” is a brief non-normative synopsis of the semantics.
- The OCL is defined using the OCL syntax defined in UML 1.3. The “context” clause names an “abstract” operation or method on an M1 level interface whose semantics is specified. The name of the real operation(s) or method(s) will depend on the mapping. The semantics are expressed as post-conditions for these methods.

3.9.2.3 Notation for Helper Functions.

OCL Helper Functions are described in the following notation:

```
[S-xxx]      HelperName
description:  brief english description

context ClassifierName::HelperName(...) : ...
post: result = ...
```

The meaning of the above is as follows:

- “[S-xxx]” is the cross reference tag for the helper function that may be used elsewhere in this document.
- “HelperName” is the name of the Helper Function.
- The “description” is a brief non-normative synopsis of the Helper’s semantics.
- The OCL for the Helper is defined using the OCL syntax defined in UML 1.3. The “context” clause names a notional helper function on a ModelElement whose semantic then specified. These notional functions are not intended to be callable by client code.

3.9.3 OCL Usage in the MOF Model specification

The OCL language was designed as a part of the UML specification. As such, the OCL semantics are specified in terms of UML concepts and constructs. Some of these concepts do not match MOF concepts exactly. Accordingly, it is necessary to reinterpret parts of the OCL specification so that it can be used in MOF Model’s Constraints and other semantics aspects of the MOF Model.

3.9.3.1 UML AssociationEnds versus MOF References

In the UML version of OCL, the dot (“.”) and arrow (“->”) operators are used to access Attribute values, and to navigate Associations. Consider an OCL expression of the form:

<expr> "." <identifier>

Assuming that "<expr>" evaluates to an object, the value of the expression is either the value of an Attribute named "<identifier>" for the object or another object obtained by navigating a link in a binary Association which has "<identifier>" as an Association End name.

In this context (i.e., the definition of the MOF Model), the "<identifier>" is interpreted differently. In the MOF Model, the interfaces for navigating Associations are specified using References rather than AssociationEnds. Thus in the MOF version of OCL, link navigation is expressed using the name of a Reference for the "<expr>" object as the "<identifier>". However, the overall meaning is analogous to the UML case.

3.9.3.2 *Helper functions are not MOF Operations*

In the UML version of OCL, object behavior is invoked by an expression of the form:

<expr> "." <identifier> "(" ... ")"

where "<identifier>" names a UML Operation or Method on the object obtained by evaluating "<expr>."

In the MOF Model specification, the above expression invokes behavior defined by either a MOF Operation, or a helper function. The distinction between conventional UML and its usage here is that helper functions have no defined connection with any internal or external interfaces in a MOF Model server. Indeed, they need not exist at all as implementation artifacts.

3.9.3.3 *Post-conditions on MOF Model objects*

Rules Chapter [C-2], Chapter [C-3] and Chapter [C-4] are intended to define post-conditions on all operations on ModelElement objects. This is expressed in the MOF Model OCL by giving a Class rather than an Operation as the "context" for the OCL rules. It is not clear that this is allowed by UML OCL.

3.9.3.4 *OCL evaluation order*

The UML OCL specification does not define an evaluation order for OCL expressions in general, and for boolean operators in particular. This is OK when OCL is used as an abstract specification language, as it is in the UML specification. However it causes problems when OCL expressions may be directly evaluated. These problems arise in OCL that traverses cyclic graphs (e.g. Chapter [O-1]) or raises exceptions (e.g. Chapter [S-6])

The MOF Model semantic specification touches on some of these issues; e.g. when traversing an cyclical Imports graph. Therefore, the MOF Model usage of OCL makes the following assumptions about OCL expression evaluation order:

- In general, a MOF OCL expression is assumed to be evaluated by evaluating its sub-expressions in order, starting with the leftmost sub-expression and ending with the rightmost. The sub-expressions are delimited according to the standard OCL operator precedence rules. If evaluation of one of the sub-expressions raises an exception, the remaining sub-expressions are not evaluated.
- The above does not apply to the boolean operators “**and**,” “**or**,” “**implies**,” and “**if-then-else**.” These are evaluated with short-circuiting as follows:
 - In the expression “<expr1> **and** <expr2>,” “<expr2>” is only evaluated if “<expr1>” evaluates to true.
 - In the expression “<expr1> **or** <expr2>,” “<expr2>” is only evaluated if “<expr1>” evaluates to false.
 - In the expression “<expr1> **implies** <expr2>,” “<expr2>” is only evaluated if “<expr1>” evaluates to true.
 - In the expression “**if** <expr1> **then** <expr2> **else** <expr3> **endif**,” “<expr2>” is only evaluated if “<expr1>” evaluates to true, and “<expr3>” is only evaluated if “<expr1>” evaluates to false.

3.9.3.5 “*OclType::allInstances*”

In UML OCL, the `type.allInstances()` is defined to return:

“The set of all instances of *type* and all of its subtypes in existence at the moment in time that the expression is evaluated”.

In the MOF Model OCL, this expression is used to refer to the set of all instances that exist within a given outermost Package extent. (Any OCL expression that required the enumeration of all instances in existence “anywhere” would be problematical, since a MOF repository does not exist in a closed world.)

3.9.3.6 “*OclType::references*”

The MOF Model OCL in rule Chapter [C-4] assumes that the signature of `OclType` (as defined in the UML OCL specification) is extended to include an operation called “references”. This is assumed to behave like the “attributes” operation, except that it returns the names of an (M3-level) Classes’ References.

3.9.3.7 *Foreign types and operations*

Some of the MOF Model OCL rules makes use of types and operations that are not predefined in OCL, not defined as Operations in the MOF Model, and not defined as Helper functions. Examples include:

- Some rules use operations on CORBA’s built-in Any and TypeCode data types. Indeed one rule uses a constructor for TypeCodes. In each case, the intended meaning should be self-evident.
- Some rules use of the TypeKind enumeration type. The meaning should be self evident.

- Rule Chapter [C-3] makes use of the CORBA Object::non_existent operation to assert that an object must continue to exist.
- Rules Chapter [C-2] and Chapter [C-4] use operations defined in the RefObject and RefBaseObject interfaces to access the meta-objects that represent the MOF Model. It should be understood that this is not intended to imply that a MOF Model server is required to make these objects available at runtime.

3.9.4 The MOF Model Constraints

[C-1] MustBeContainedUnlessPackage
format1: MUST_BE_CONTAINED_UNLESS_PACKAGE
format2: must_be_contained_unless_package
evaluation policy: deferred
description: A ModelElement that is not a Package must have a container.
context ModelElement
inv:
 not self.oclIsTypeOf(Package) **implies**
 self.container -> size = 1

[C-2] FrozenAttributesCannotBeChanged
format1: FROZEN_ATTRIBUTES_CANNOT_BE_CHANGED
format2: frozen_attributes_cannot_be_changed
evaluation policy: immediate
description: The attribute values of a ModelElement which is frozen cannot be changed.
context ModelElement
inv:
 self.isFrozen() implies
 let myTypes = **self**.oclType() -> allSupertypes() ->
 includes(**self**.oclType()) **in**
 let myAttrs : Set(Attribute) =
 self.RefBaseObject::refMetaObject() ->
 asOclType(Class) ->
 findElementsByTypeExtended(Attribute) **in**
 myAttrs -> forAll(a |
 self.RefObject::refValue@pre(a) =
 self.RefObject::refValue(a))

[C-3] FrozenElementsCannotBeDeleted
 format1: FROZEN_ELEMENTS_CANNOT_BE_DELETED
 format2: frozen_elements_cannot_be_deleted
 evaluation policy: immediate
 description: A frozen ModelElement which is in a frozen Namespace can only be deleted, by deleting the Namespace.

```
context ModelElement
post:
  (self.isFrozen@pre() and
   self.container@pre -> notEmpty and
   self.container.isFrozen@pre()) implies
  (self.container.Object::non_existent() or
   not self.Object::non_existent())
```

[C-4] FrozenDependenciesCannotBeChanged
 format1: FROZEN_DEPENDENCIES_CANNOT_BE_CHANGED
 format2: frozen_dependencies_cannot_be_changed
 evaluation policy: immediate
 description: The link sets that express dependencies of a frozen Element on other Elements cannot be explicitly changed.

```
context ModelElement
post:
  self.isFrozen() implies
  let myClasses = self.oclType() -> allSupertypes() ->
    includes(self.oclType()) in
  let myRefs = Set(Reference) =
    self.RefBaseObject::refMetaObject() ->
    asOclType(Class) ->
    findElementsByTypeExtended(Reference) in
  let myDepRefs = myRefs ->
    select(r |
      Set{"contents", "constraints", "supertypes",
        "type", "referencedEnd", "exceptions",
        "importedNamespace", "elements"} ->
      includes(r.name)) in
  myDepRefs ->
    forAll(r |
      self.RefObject::refValue@pre(r) =
        self.RefObject::refValue(r))
```

- [C-5] ContentNamesMustNotCollide
format1: CONTENT_NAMES_MUST_NOT_COLLIDE
format2: content_names_must_not_collide
evaluation policy: immediate
description: The names of the contents of a Namespace must not collide.
context Namespace
inv: `self.contents.forAll(
e1, e2 | e1.name = e2.name implies r1 = r2)`
- [C-6] SupertypeMustNotBeSelf
format1: SUPERTYPE_MUST_NOT_BE_SELF
format2: supertype_must_not_be_self
evaluation policy: immediate
description: A Generalizable Element cannot be its own direct or indirect supertype.
context GeneralizableElement
inv: `self.allSupertypes() -> forAll(s | s <> self)`
- [C-7] SupertypeKindMustBeSame
format1: SUPERTYPE_KIND_MUST_BE_SAME
format2: supertype_kind_must_be_same
evaluation policy: immediate
description: A supertypes of a GeneralizableElement must be of the same kind as the GeneralizableElement itself.
context GeneralizableElement
inv: `self.supertypes -> forAll(s | s.oclcType() = self.oclcType())`
- [C-8] ContentsMustNotCollideWithSupertypes
format1: CONTENTS_MUST_NOT_COLLIDE_WITH_SUPERTYPES
format2: contents_must_not_collide_with_supertypes
evaluation policy: immediate
description: The names of the contents of a GeneralizableElement should not collide with the names of the contents of any direct or indirect supertype.
context GeneralizableElement
inv:
`let superContents = self.allSupertypes() ->
collect(s | s.contents) in
self.contents ->
forAll(m1 |`

```

superContents ->
  forAll(m2 |
    m1.name = m2.name implies m1 = m2))

```

[C-9] DiamondRuleMustBeObeyed
format1: DIAMOND_RULE_MUST_BE_OBEYED
format2: diamond_rule_must_be_obeyed
evaluation policy: immediate
description: Multiple inheritance must obey the “Diamond Rule”.
context GeneralizableElement
inv:
 let superNamespaces =
 self.supertypes -> collect(s | s.extendedNamespace) **in**
 superNamespaces -> asSet -> isUnique(s | s.name)

[C-10] NoSupertypesAllowedForRoot
format1: NO_SUPERTYPES_ALLOWED_FOR_ROOT
format2: no_supertypes_allowed_for_root
evaluation policy: immediate
description: If a Generalizable Element is marked as a “root”, it cannot have any supertypes.
context GeneralizableElement
inv: **self**.isRoot **implies self**.supertypes -> isEmpty

[C-11] SupertypesMustBeVisible
format1: SUPERTYPES_MUST_BE_VISIBLE
format2: supertypes_must_be_visible
evaluation policy: deferred
description: A GeneralizableElement’s immediate supertypes must all be visible to it.
context GeneralizableElement
inv: **self**.supertypes -> forAll(s | **self**.isVisible(s))

[C-12] NoSubtypesAllowedForLeaf
format1: NO_SUBTYPES_ALLOWED_FOR_LEAF
format2: no_subtypes_allowed_for_leaf
evaluation policy: immediate
description: A GeneralizableElement cannot inherit from a GeneralizableElement defined as a “leaf”.
context GeneralizableElement
inv: **self**.supertypes -> forAll(s | **not** s.isLeaf)

- [C-13] AssociationsCannotBeTypes
format1: ASSOCIATIONS_CANNOT_BE_TYPES
format2: associations_cannot_be_types
evaluation policy: immediate
description: An Association cannot be the type of a TypedElement.
context TypedElement
inv: **not** **self.type.oclIsKindOf**(Association)
- [C-14] TypeMustBeVisible
format1: TYPE_MUST_BE_VISIBLE
format2: type_must_be_visible
evaluation policy: deferred
description: A TypedElement can only have a type that is visible to it.
context TypedElement
inv: **self.isVisible**(**self.type**)
- [C-15] ClassContainmentRules
format1: CLASS_CONTAINMENT_RULES
format2: class_containment_rules
evaluation policy: immediate
description: A Class may contain only Classes, DataTypes, Attributes, References, Operations, Exceptions, Constraints and Tags.
context Class
inv:
 Set{Class, DataType, Attribute, Reference, Operation,
 Exception, Constraint, Tag} ->
 includesAll(**self.contentTypes**())
- [C-16] AbstractClassesCannotBeSingleton
format1: ABSTRACT_CLASSES_CANNOT_BE_SINGLETON
format2: abstract_classes_cannot_be_singleton
evaluation policy: deferred
description: A Class that is marked as abstract cannot also be marked as singleton.
context Class
inv: **self.isAbstract** **implies not** **self.isSingleton**

[C-17] DataTypeContainmentRules
 format1: DATA_TYPE_CONTAINMENT_RULES
 format2: data_type_containment_rules
 evaluation policy: immediate
 description: A DataType may contain only TypeAliases, Constraints and Tags.

context DataType
inv:
 Set{TypeAlias, Constraint, Tag} ->
 includesAll(**self**.contentTypes())

[C-18] ThisTypecodeNotSupported
 format1: THIS_TYPECODE_NOT_SUPPORTED
 format2: this_typecode_not_supported
 evaluation policy: deferred
 description: The typeCode of a DataType must denote a CORBA 2.2 compliant object type or data type.

context DataType
inv:
 self.typeCode.allTypeKinds() ->
 excludes(Set{#tk_void, #tk_Principal, #tk_null, #tk_except,
 #tk_value, #tk_value_box, #tk_native,
 #tk_abstract_interface})

[C-19] DataTypesHaveNoSupertypes
 format1: DATA_TYPES_HAVE_NO_SUPERTYPES
 format2: data_types_have_no_supertypes
 evaluation policy: immediate
 description: Inheritance / generalization is not applicable to DataTypes.

context DataType
inv: **self**.supertypes -> isEmpty

[C-20] DataTypesCannotBeAbstract
 format1: DATA_TYPES_CANNOT_BE_ABSTRACT
 format2: data_types_cannot_be_abstract
 evaluation policy: immediate
 description: A DataType cannot be abstract.

context DataType
inv: **not self.isAbstract**

- [C-21] ReferenceMultiplicityMustMatchEnd
 format1: REFERENCE_MULTPLICITY_MUST_MATCH_END
 format2: reference_multiplicity_must_match_end
 evaluation policy: deferred
 description: The multiplicity for a Reference must be the same as the multiplicity for the referenced AssociationEnd.
context Reference
inv: **self**.multiplicity = **self**.referencedEnd.multiplicity
- [C-22] ReferenceMustBeInstanceScoped
 format1: REFERENCE_MUST_BE_INSTANCE_SCOPED
 format2: reference_must_be_instance_scoped
 evaluation policy: immediate
 description: Classifier scoped References are not meaningful in the current M1 level computational model.
context Reference
inv: **self**.scope = #instance_level
- [C-23] ChangeableReferenceMustHaveChangeableEnd
 format1: CHANGEABLE_REFERENCE_MUST_HAVE_CHANGEABLE_END
 format2: changeable_reference_must_have_changeable_end
 evaluation policy: deferred
 description: A Reference can be changeable only if the referenced AssociationEnd is also changeable.
context Reference
inv: **self**.isChangeable = **self**.referencedEnd.isChangeable
- [C-24] ReferenceTypeMustMatchEndType
 format1: REFERENCE_TYPE_MUST_MATCH_END_TYPE
 format2: reference_type_must_match_end_type
 evaluation policy: deferred
 description: The type attribute of a Reference and its referenced AssociationEnd must be the same.
context Reference
inv: **self**.type = **self**.referencedEnd.type

[C-25] ReferencedEndMustBeNavigable
 format1: REFERENCED_END_MUST_BE_NAVIGABLE
 format2: referenced_end_must_be_navigable
 evaluation policy: deferred
 description: A Reference is only allowed for a navigable AssociationEnd
context Reference
inv: **self**.referencedEnd.isNavigable

[C-26] ContainerMustMatchExposedType
 format1: CONTAINER_MUST_MATCH_EXPOSED_TYPE
 format2: container_must_match_exposed_type
 evaluation policy: deferred
 description: The containing Class for a Reference must be equal to or a subtype of
 the type of the Reference's exposed AssociationEnd.
context Reference
inv:
 self.container.allSupertypes() -> including(**self**) ->
 includes(**self**.referencedEnd.otherEnd.type)

[C-27] ReferencedEndMustBeVisible
 format1: REFERENCED_END_MUST_BE_VISIBLE
 format2: referenced_end_must_be_visible
 evaluation policy: deferred
 description: The referenced AssociationEnd for a Reference must be visible from
 the Reference.
context Reference
inv: **self**.isVisible(**self**.referencedEnd)

[C-28] OperationContainmentRules
 format1: OPERATION_CONTAINMENT_RULES
 format2: operation_containment_rules
 evaluation policy: immediate
 description: An Operation may only contain Parameters, Constraints and Tags.
context Operation
inv:
 Set{Parameter, Constraint, Tag} ->
 includesAll(**self**.contentTypes())

[C-29] OperationsHaveAtMostOneReturn
 format1: OPERATIONS_HAVE_AT_MOST_ONE_RETURN
 format2: operations_have_at_most_one_return
 evaluation policy: immediate
 description: An Operation may have at most one Parameter whose direction is
 "return".

context Operation

inv:

```

self.contents ->
  select(c | c.oclIsTypeOf(Parameter)) ->
    select(p : Parameter | p.direction = #return_dir) ->
      size < 2

```

[C-30] OperationExceptionsMustBeVisible
 format1: OPERATION_EXCEPTIONS_MUST_BE_VISIBLE
 format2: operation_exceptions_must_be_visible
 evaluation policy: deferred
 description: The Exceptions raised by an Operation must be visible to the
 Operation.

context Operation

inv: **self**.exceptions -> forAll(e | **self**.isVisible(e))

[C-31] ExceptionContainmentRules
 format1: EXCEPTION_CONTAINMENT_RULES
 format2: exception_containment_rules
 evaluation policy: immediate
 description: An Exception may only contain Parameters and Tags.

context Exception

inv: Set{Parameter, Tag}) -> includesAll(**self**.contentTypes())

[C-32] ExceptionsHaveOnlyOutParameters
 format1: EXCEPTIONS_HAVE_ONLY_OUT_PARAMETERS
 format2: exceptions_have_only_out_parameters
 evaluation policy: immediate
 description: An Exception's Parameters must all have the direction "out".

context Exception

inv:

```

self.contents ->
  select(c | c.oclIsTypeOf(Parameter)) ->
    forAll(p : Parameter | p.direction = #out_dir)

```

[C-33] AssociationContainmentRules
format1: ASSOCIATIONS_CONTAINMENT_RULES
format2: associations_containment_rules
evaluation policy: immediate
description: An Association may only contain AssociationEnds, Constraints and Tags.

context Association

inv:

```
Set{AssociationEnd, Constraint, Tag} ->  
includesAll(self.contentTypes())
```

[C-34] AssociationsHaveNoSupertypes
format1: ASSOCIATIONS_HAVE_NO_SUPERTYPES
format2: associations_have_no_supertypes
evaluation policy: immediate
description: Inheritance / generalization is not applicable to Associations.

context Association

inv: self.supertypes -> isEmpty

[C-35] AssociationMustBeRootAndLeaf
format1: ASSOCIATIONS_MUST_BE_ROOT_AND_LEAF
format2: associations_must_be_root_and_leaf
evaluation policy: immediate
description: The values for “isLeaf” and “isRoot” on an Association must be true.

context Association

inv: self.isRoot and self.isLeaf

[C-36] AssociationsCannotBeAbstract
format1: ASSOCIATIONS_CANNOT_BE_ABSTRACT
format2: associations_cannot_be_abstract
evaluation policy: immediate
description: An Association cannot be abstract.

context Association

inv: not self.isAbstract

[C-37] AssociationsMustBePublic
 format1: ASSOCIATIONS_MUST_BE_PUBLIC
 format2: associations_must_be_public
 evaluation policy: immediate
 description: Associations must have visibility of “public”.
context Association
inv: **self**.visibility = #public_vis

[C-38] AssociationsMustBeBinary
 format1: ASSOCIATIONS_MUST_BE_BINARY
 format2: associations_must_be_binary
 evaluation policy: immediate
 description: An Association must be binary; i.e. it must have exactly two AssociationEnds.
context Association
inv: **self**.contents ->
 select(c | c.oclIsTypeOf(AssociationEnd)) -> size = 2

[C-39] EndTypeMustBeClass
 format1: END_TYPE_MUST_BE_CLASS
 format2: end_type_must_be_class
 evaluation policy: immediate
 description: The type of an AssociationEnd must be Class.
context AssociationEnd
inv: **self**.type.oclIsTypeOf(Class)

[C-40] EndsMustBeUnique
 format1: ENDS_MUST_BE_UNIQUE
 format2: ends_must_be_unique
 evaluation policy: immediate
 description: The “isUnique” flag in an AssociationEnd’s multiplicity must be true.
context AssociationEnd
inv:
 (**self**.multiplicity.upper > 1 or
 self.multiplicity.upper = UNBOUNDED) **implies**
 self.multiplicity.isUnique

[C-41] CannotHaveTwoOrderedEnds
format1: CANNOT_HAVE_TWO_ORDERED_ENDS
format2: cannot_have_two_ordered_ends
evaluation policy: deferred
description: An Association cannot have two AssociationEnds marked as
 “ordered”.

context AssociationEnd
inv:
 self.multiplicity.isOrdered **implies**
 not self.otherEnd.multiplicity.isOrdered

[C-42] CannotHaveTwoAggregateEnds
format1: CANNOT_HAVE_TWO_AGGREGATE_ENDS
format2: cannot_have_two_aggregate_ends
evaluation policy: deferred
description: An Association cannot have an aggregation semantic specified for
 both AssociationEnds.

context AssociationEnd
inv:
 self.aggregation <> #none **implies self**.otherEnd = #none

[C-43] PackageContainmentRules
format1: PACKAGE_CONTAINMENT_RULES
format2: package_containment_rules
evaluation policy: immediate
description: A Package may only contain Packages, Classes, DataTypes,
 Associations, Exceptions, Constraints, Imports and Tags.

context Package
inv:
 Set{Package, Class, DataType, Association, Exception,
 Constraint, Import, Tag}) -> includesAll(**self**.contentTypes)

[C-44] PackagesCannotBeAbstract
format1: PACKAGES_CANNOT_BE_ABSTRACT
format2: packages_cannot_be_abstract
evaluation policy: immediate
description: Packages cannot be declared as abstract.

context Package
inv: **not self**.isAbstract

[C-45] ImportedNamespaceMustBeVisible
 format1: IMPORTED_NAMESPACE_MUST_BE_VISIBLE
 format2: imported_namespace_must_be_visible
 evaluation policy: deferred
 description: The Namespace imported by an Import must be visible to the Import's containing Package.

context Import
inv: `self.container.isVisible(self.importedNamespace)`

[C-46] CanOnlyImportPackagesAndClasses
 format1: CAN_ONLY_IMPORT_PACKAGES_AND_CLASSES
 format2: can_only_import_packages_and_classes
 evaluation policy: immediate
 description: It is only legal for a Package to import or cluster Packages or Classes.

context Import
inv:
 `self.imported.oclIsTypeOf(Class) or`
 `self.imported.oclIsTypeOf(Package)`

[C-47] CannotImportSelf
 format1: CANNOT_IMPORT_SELF
 format2: cannot_import_self
 evaluation policy: deferred
 description: Packages cannot import or cluster themselves.

context Import
inv: `self.container <> self.imported`

[C-48] CannotImportNestedComponents
 format1: CANNOT_IMPORT_NESTED_COMPONENTS
 format2: cannot_import_nested_components
 evaluation policy: deferred
 description: Packages cannot import or cluster Packages or Classes that they contain.

context Import
inv: `not self.container.allContents() -> includes(self.imported)`

[C-49] NestedPackagesCannotImport
format1: NESTED_PACKAGES_CANNOT_IMPORT
format2: nested_packages_cannot_import
evaluation policy: deferred
description: Nested Packages cannot import or cluster other Packages or Classes.
context Import
inv:
 self.container -> notEmpty **implies**
 self.container -> asSequence -> first -> container -> isEmpty

[C-50] CannotConstrainThisElement
format1: CANNOT_CONSTRAIN_THIS_ELEMENT
format2: cannot_constrain_this_element
evaluation policy: immediate
description: Constraints, Tags, Imports, TypeAliases and Constants cannot be
 constrained.
context Constraint
inv:
 self.constrainedElements ->
 forall(c |
 not Set{Constraint, Tag, Imports,
 TypeAlias, Constant} ->
 includes(c.oclType()))

[C-51] ConstraintsLimitedToContainer
format1: CONSTRAINTS_LIMITED_TO_CONTAINER
format2: constraints_limited_to_container
evaluation policy: deferred
description: A Constraint can only constrain ModelElements that are defined by or
 inherited by its immediate container.
context Constraint
inv:
 self.constrainedElements ->
 forall(c | **self**.container.extendedNamespace() ->
 includes(c))

[C-52] ConstantsValueMustMatchType
format1: CONSTANTS_VALUE_MUST_MATCH_TYPE
format2: constants_value_must_match_type
evaluation policy: deferred
description: The type of a Constant and the type of its value must be the same.
context Constant
inv: `self.value.type -> equals(self.type -> mapToTypecode())`

[C-53] ConstantsTypeMustBeSimpleDataType
format1: CONSTANTS_TYPE_MUST_BE_SIMPLE_DATA_TYPE
format2: constants_type_must_be_simple_data_type
evaluation policy: immediate
description: The type of a Constant must be a CORBA data type that is legal for a CORBA 2.3 constant declaration.
context Constant
inv:
`self.type.ocIsKindOf(DataType) and`
`Set{#tk_short, #tk_ushort, #tk_long, #tk_ulong, #tk_char,`
`#tk_octet, #tk_float, #tk_double, #tk_boolean, #tk_string,`
`#tk_wchar, #tk_wstring, #tk_longlong, #tk_ulonglong,`
`#tk_longdouble, #tk_fixed} ->`
`includes(self.type.asType(DataType).typecode.`
`unwindAliases().kind)`

[C-54] LowerCannotBeNegativeOrUnbounded
format1: LOWER_CANNOT_BE_NEGATIVE_OR_UNBOUNDED
format2: lower_cannot_be_negative_or_unbounded
evaluation policy: immediate
description: The “lower” bound of an MultiplicityType to be “Unbounded”.
context MultiplicityType
inv: `self.lower >= 0 and self.lower <> Unbounded`

[C-55] LowerCannotExceedUpper
format1: LOWER_CANNOT_EXCEED_UPPER
format2: lower_cannot_exceed_upper
evaluation policy: immediate
description: The “lower” bound of a MultiplicityType cannot exceed the “upper” bound.
context MultiplicityType
inv: `self.lower <= self.upper or self.upper = Unbounded`

[C-56] UpperMustBePositive
 format1: UPPER_MUST_BE_POSITIVE
 format2: upper_must_be_positive
 evaluation policy: immediate
 description: The “upper” bound of a MultiplicityType cannot be less than 1.
context MultiplicityType
inv: `self.upper >= 1 or self.upper = Unbounded`

[C-57] MustBeUnorderedNonunique
 format1: MUST_BE_UNORDERED_NONUNIQUE
 format2: must_be_unordered_nonunique
 evaluation policy: immediate
 description: If a MultiplicityType specifies bounds of [0..1] or [1..1] the “is_ordered” and “is_unique” values must be false.
context MultiplicityType
inv:
 `self.upper = 1 implies`
 `(not self.isOrdered and not self.isUnique)`

3.9.5 *Semantic specifications for some Operations, derived Attributes and Derived Associations*

[S-1] allSupertypes
 kind: query Operation
 description: The value is the closure of the ‘Generalizes’ Association from the perspective of a subtype. Note that the sequence of all supertypes has a well defined order.
context GeneralizableElement::allSupertypes() :
 Sequence(GeneralizableElement)
post: `result = self.allSupertypes2(Set{})`

[S-2] otherEnd
 kind: query Operation
 description: The value of is the other AssociationEnd for this Association.
context AssociationEnd::otherEnd() : AssociationEnd
post: `result = self.container.contents ->`
 `select(c | c.oclIsKindOf(AssociationEnd) and c <> self)`

[S-3] isVisible
 kind: query Operation
 description: Determines whether or not “otherElement” is visible for the definition of this element. (NB: As an interim measure, the OCL states that everything is visible!)
context ModelElement::isVisible(
 otherElement : ModelElement): boolean
post: result = true

[S-4] findRequiredElements
 kind: query Operation
 description: Selects a subset of a ModelElements immediate or recursive dependents.
context ModelElement::isRequiredBecause(
 kinds : Sequence(DependencyKind),
 recursive : boolean) : Sequence(ModelElement)
post: result =
 if kinds -> includes(“all”)
 then
 self.findRequiredElements(
 Set{“constraint”, “container”, “constrained elements”,
 “specialization”, “import”, “contents”, “signature”,
 “tagged elements”, “type definition”,
 “referenced ends”})
 else
 if recursive
 then
 self.recursiveFindDeps(kinds, Set{**self**})
 else
 kinds -> collect(k : **self**.findDepsOfKind(k)) -> asSet()
 endif
 endif

[S-5] isRequiredBecause
 kind: query Operation
 description: Returns the DependencyKind that describes the dependency between this element and “other”.
context ModelElement::findRequiredElements(
 other : ModelElement,
 reason : **out** DependencyKind) : boolean
post: -- NB, if there is more than one dependency between self
 -- and ‘other’, the selection of the ‘reason’ is defined
 -- to be non-deterministic ... not deterministic as a
 -- left to right evaluation of the OCL implies.
 reason = (
 if self -> isDepOfKind(“constraint”, other)
 then
 “constraint”
)

```
else
  if self -> isDepOfKind("container", other)
  then
    "container"
  else
    if self -> isDepOfKind("constrained elements", other)
    then
      "constrained elements"
    else
      if self -> isDepOfKind("specialization", other)
      then
        "specialization"
      else
        if self -> isDepOfKind("import", other)
        then
          "import"
        else
          if self -> isDepOfKind("contents", other)
          then
            "contents"
          else
            if self -> isDepOfKind("signature", other)
            then
              "signature"
            else
              if self -> isDepOfKind("tagged elements", other)
              then
                "tagged elements"
              else
                if self -> isDepOfKind("type definition", other)
                then
                  "type definition"
                else
                  if self -> isDepOfKind("referenced ends", other)
                  then
                    "referenced ends"
                  else
                    if self -> dependsOn() -> notEmpty()
                    then
                      "indirect"
                    else
                      ""
                    endif
                  endif
                endif
              endif
            endif
          endif
        endif
      endif
    endif
  endif) and
result = (reason <> "")
```

[S-6] lookupElement
 kind: query Operation
 description: Returns the ModelElement in the Namespace whose name is equal to
 "name", or raises an exception.

```

context Namespace::lookupElement(name : NameType) : ModelElement
post: result =
  let elems = self.contents -> select(m | m.name = name) in
  if elems -> size = 0
  then
    -- Raise exception NameNotFound
  else
    elems -> first -- should only be one
  endif

```

[S-7] resolveQualifiedName
 kind: query Operation
 description: Returns the ModelElement that "qualifiedName" resolves to or raises
 an exception

```

context Namespace::resolveQualifiedName(
  qualifiedName : Sequence(NameType)) : ModelElement
pre: qualifiedName -> size >= 1
post: result =
  let elems = self.contents ->
    select(m | m.name = qualifiedName -> first) in
  if elems -> size = 0
  then
    -- Raise exception NameNotResolved ...
  else
    if qualifiedName -> size = 1
    then
      elems -> first -- there should only be one
    else
      if not elems -> first -> oclIsOfKind(Namespace)
      then
        -- Raise exception NameNotResolved ...
      else
        let rest = qualifiedName ->
          subSequence(2, qualifiedName -> size) in
          elems -> first -> resolveQualifiedName(rest)
        endif
      endif
    endif
  endif

```

[S-8] nameIsValid
 kind: query Operation
 description: Returns true if “proposedName” is a valid name that could be used for a new containedElement of this Namespace.

```

context Namespace::nameIsValid(
    proposedName : NameType) : boolean
post: result =
    self.extendedNamespace ->
        forAll(e | not e.name = proposedName)
  
```

[S-9] findElementsByType
 kind: query Operation
 description: Returns a subset of the contained elements. If “includeSubtypes” is false, the result consists of instances of “ofType”. If it is true, instances of subClasses are included.

```

context Namespace::findElementsByType(
    ofType : Class,
    includeSubtypes : boolean) : Sequence(ModelElement)
post: result =
    if includeSubtypes
    then
        self.contents -> select(m | m.oclIsOfKind(ofType))
    else
        self.contents -> select(m | m.oclIsOfType(ofType))
    endif
  
```

[S-10] lookupElementExtended
 kind: query Operation
 description: Returns the ModelElement whose name is equal to “name” in the extended namespace of this GeneralizableElement, or raises an exception.

```

context Namespace::lookupElementExtended(
    name : NameType) : ModelElement
post: result =
    let elems = self -> extendedNamespace ->
        select(m | m.name = name) in
    if elems -> size = 0
    then
        -- Raise exception NameNotFound
    else
        elems -> first -- should only be one
    endif
  
```

[S-11] findElementsByTypeExtended
 kind: query Operation
 description: Returns a subset of the contained, inherited or imported elements. If “includeSubtypes” is false, the result consists of instances of “ofType”. If it is true, instances of subClasses are included.

```

context GeneralizeableElement::findElementsByTypeExtended(
    ofType : Class,
    includeSubtypes : boolean) : Sequence(ModelElement)
post: result =
  if includeSubtypes
  then
    self.extendedNamespace -> select(m | m.oclIsOfKind(ofType))
  else
    self.extendedNamespace -> select(m | m.oclIsOfType(ofType))
  endif

```

[S-12] qualifiedName
 kind: readonly derived Attribute
 description: The qualified name gives the sequence of names of the containers of this ModelElement starting with the outermost.

```

context ModelElement::qualifiedName() : Sequence(ModelElement)
post: result =
  if self.container -> notEmpty
  then
    self.container.qualifiedName() -> append(self.name)
  else
    self.name
  endif

```

[S-13] Exposes
 kind: derived Association
 description: This association relates a Reference to the exposed AssociationEnd of an Association that corresponds to its referencedEnd.

```

context Reference
inv: AssociationEnd.allInstances ->
  forAll(
    a |
    self.references = a implies self.exposes = a.otherEnd and
    not self.references = a implies self.exposes <> a.otherEnd)

```

[S-14] DependsOn
 kind: derived Association
 description: This association relates a ModelElement to the other ModelElements whose definition it depends on.

```

context ModelElement
inv: self.findRequiredElements("all", true)

```

3.9.6 OCL Helper functions

[O-1] allSupertypes2

description: Helper function for the allSupertypes operation.

```

context GeneralizableElement::allSupertypes2(
    visited : Set(GeneralizableElement)) :
    Sequence(GeneralizableElement)
post: result =
    if (visited -> includes(self))
    then
        Sequence{}
    else
        let mySupers : Sequence(GeneralizableElement) =
            self.supertypes ->
                collect(s |
                    s.allSupertypes2(visited ->
                        including(self))) in
            mySupers ->
                iterate(s2 : GeneralizableElement;
                    a : Sequence(GeneralizableElement) = Sequence{} |
                    if a -> includes(s2)
                    then
                        a
                    else
                        a -> append(s2)
                    endif)
    
```

[O-2] allTypeKinds

description: Return the TypeCode's kind and the kinds of embedded TypeCodes

```

context TypeCode::allTypeKinds() : Set(TCKind)
post: result =
    let memberKinds = self.memberTypecodes ->
        collect(mt | mt.allTypeKinds()) in
    let contentKinds =
        if Set{tk_sequence, tk_array, tk_alias} ->
            includes(self.kind())
        then
            self.content_type().allTypeKinds()
        else
            Set{}
        endif in
    let discrimKinds =
        if self.kind() = tk_union
        then
            Set{self.discriminator_type().kind}
        else
            Set{}
        endif in
    Set{self.kind()} -> union(memberKinds) ->
        union(contentKinds) -> union(discrimKinds)
    
```


[O-3] memberTypecodes
description: returns the set (possibly empty) of member TypeCodes for a TypeCode

```

context TypeCode::memberTypecodes() : Set(TypeCode)
post: result =
  if Set{tk_struct, tk_union, tk_except} ->
    includes(self.kind())
  then
    Set{0..self.member_count()} ->
      collect(i | self.member_type(i).allTypeKinds())
  else
    Set{}
  endif

```

[O-4] unwindAliases
description: dealiases a TypeCode

```

context TypeCode::unwindAliases() : TypeCode
post: result =
  if self.kind() = tk_alias then
    self.content_type().unwindAliases()
  else
    self
  endif

```

[O-5] extendedNamespace
description: The extendedNamespace of a Namespace is its contents, the contents of its supertypes and any Namespaces that it imports.

```

context Namespace::extendedNamespace() : Set(ModelElement)
post: result =
  self.contents

context GeneralizableElement::extendedNamespace : Set(ModelElement)
post: result =
  self.contents ->
    union(self.allSupertypes() -> collect(s | s.contents))

context Package::extendedNamespace : Set(ModelElement)
post: result =
  let ens = self.contents ->
    union(self.allSupertypes() -> collect(s | s.contents)) in
  let imports = ens -> select(e | e.oclKindOf(Import)) ->
    collect(i : Import | i.imported) in
  ens -> union(imports)

```

[O-6] contentTypes
description: The set of OCL types for a Namespace's contents.

```

context Namespace::contentTypes() : Set(OCLType)
post: result = self.contents -> collect(m | m.oclType()) -> asSet

```

[O-7] mapToTypecode

description: The typecode corresponding to a Classifier.

context Classifier::mapToTypecode() : TypeCode

```

post: result =
  if self.oclIsTypeOf(DataType)
  then
    self.typecode
  else
    new TypeCode(tk_objref,
                  self -> format1Name,
                  self -> repositoryId)

```

[O-8] format1Name

description: The simple name of the element converted to words and reassembled according to the “format1” rules; see “IDL Identifier Format 1” on page 5-39.

context ModelElement::format1Name() : string

post: result = ...

[O-9] repositoryId

description: The qualified name of the element converted into a standard CORBA repositoryId string.

context ModelElement::repositoryId() : string

post: result = ...

[O-10] recursiveFindDeps

description: The set of ModelElements which recursively depend on this one.

context ModelElement::recursiveFindDeps(
 kinds : Sequence(DependencyKind),
 seen : Set(ModelElement)) : Set(ModelElement)

```

post: result =
  let seen2 = seen ->
    collect(m | kinds ->
      collect(k | m.findDepsOfKind(k)) -> asSet) in
  if seen2 = seen
  then
    seen
  else
    self.recursiveFindDeps(kinds, seen2)
  endif

```

[O-11] isDepOfKind

description: Returns true if this element depends on 'other' with a dependency of 'kind'.

```
context ModelElement::isDepOfKind(
    kind : DependencyKind,
    other : ModelElement) : boolean
post: result = self -> findDepsOfKind(kind) -> includes(other)
```

[O-12] findDepsOfKind

description: The set of ModelElements which this one Depends on with "kind" dependency.

```
context ModelElement::findDepsOfKind(
    kind : DependencyKind) : Sequence(ModelElement)
post: result =
    if kind = "constraint"
    then
        self.constraints()
    else
        if kind = "container"
        then
            self.container()
        else
            if kind = "constrained elements" and
                self -> isOclTypeOf(Constraint)
            then
                self -> oclAsType(Constraint) -> constrainedElements()
            else
                if kind = "specialization" and
                    self -> isOclKindOf(GeneralizableElement)
                then
                    self -> oclAsType(GeneralizableElement) -> supertypes()
                else
                    if kind = "import" and self -> isOclType(Import)
                    then
                        self -> oclAsType(Import) -> importedNamespace()
                    else
                        if kind = "contents" and self -> isOclKindOf(Namespace)
                        then
                            self -> oclAsType(Namespace) -> contents()
                        else
                            if kind = "signature" and self -> isOclTypeOf(Operation)
                            then
                                self -> oclAsType(Operation) -> exceptions()
                            else
                                if kind = "tagged elements" and
                                    self -> isOclTypeOf(Tag)
                                then
                                    self -> asOclType(Tag) -> elements()
                                else
                                    if kind = "type definition" and
                                        self -> isOclKindOf(TypedElement)
                                    then
```

```
self -> asOclType(TypedElement) -> type()
if kind = "referenced ends" and
  self -> isOclKindOf(Reference)
then
  let ref = self -> asOclType(Reference) in
    ref -> referencedEnd() ->
      union(ref -> exposedEnd())
else
  Set{}
endif
endif
endif
endif
endif
endif
endif
endif
endif
```

Contents

This chapter contains the following topics.

Topic	Page
“Overview”	4-1
“MOF Values”	4-2
“Semantics of Class Instances”	4-3
“Semantics of Attributes”	4-4
“Package Composition”	4-7
“Extents”	4-9
“Semantics of Associations”	4-12
“Aggregation Semantics”	4-17
“Closure Rules”	4-18
“Recommended Copy Semantics”	4-21
“Computational Semantics”	4-22

4.1 Overview

This chapter describes the MOF’s M1-level information model, and the common principles underlying mapping specific M1-level computational models. Since it is intended to be independent of any mapping to implementation technology, the material is rather abstract.

4.2 MOF Values

A MOF meta-model is an abstract language for defining the types of meta-data. The M2-level constructs used in a meta-model map onto M1-level representations as MOF values. The types of these M1-level values can be defined using either M2-level Classes or M2-level DataTypes.

An M2-level Class defines an M1-level Instance type with the following properties:

- Instance typed objects have full object identity; i.e. it is always possible to reliably distinguish one instance (object) from another. Object identity is intrinsic and permanent, and is not dependent on other properties such as attribute values.
- Instance typed objects can be linked via an Association.

By contrast, an M2-level DataType defines a type with the following properties:

- Data typed values do *not* have full object identity; see below.
- Data typed values *cannot* be linked via an Association.

DataTypes typically fall into two groups:

- MOF data types (i.e., data types that are defined using the MOF primitive data types and constructors):
 - boolean, character types, string types, numeric types
 - enumerations, arrays, sequences and records.

By definition, MOF data types do not have object identity. Ideally, they should be available across all supported mappings with common computational semantics and inter-convertible representations.

- Native types (i.e., types that are defined in a type system that is (notionally) beyond the scope of the core MOF type system). These types will typically be specific to a given mapping, and hence should not be used where interoperability across middle-ware technologies is a goal.

Native types *may* support object identity (in some sense). However, object identity is not necessarily supported in the computational model of any given mapping.

DataTypes can also be classified as modelled types and external types. A modelled type is one whose definition is expressed within the framework of the MOF Model. By contrast, an external type is defined by some type definition mechanism outside of the MOF Model. An example of the latter is a CORBA interface type whose definition is held in a CORBA Interface Repository.

Note – The current version of the MOF specification is tied to the CORBA type system, and represents DataTypes in a CORBA specific way. Some of the more complicated CORBA data types do not have an equivalent in other type systems.

4.2.1 Semantics of Equality for MOF Values

Much of the detail of the MOF computational model depends on a notion of equality of values. For example, the precise formulation of the “no duplicates” rule for link sets depends on a definition for what it means for object type instances to be equal.

Equality of MOF Values is defined as follows:

1. Instances of Classes are equal if and only if they have the same identifier; i.e. it does *not* take into account the values of Attributes for the instances or the links involving the instances.
2. Values of MOF primitive data types are equal if and only if they have the same type and same value.
3. Values of MOF enumeration data types are equal if and only if they have the same type and the same enumerator.
4. Values of MOF array or sequence data types are equal if and only if they have the same type, the same number of members and the corresponding components are equal according to these rules.
5. Values of MOF record types are equal if and only if they have the same type and the corresponding record fields are equal according to these rules.
6. The meaning of equality of external types and native types depends on the mapping.

Note – The meaning of equality for a particular external or native type may vary depending on the mapping. For example, equality may be undefined in some mappings.

4.3 Semantics of Class Instances

An M1-level Instance is a value whose type is described by an M2-level Class. An Instance has the following properties in the MOF computational model:

- It has object identity. This has different implications depending on the mapping, but in general it means that many conceptually distinct Instance values can exist whose component values are the same.
- It has a definite lifetime. An Instance value is created in response to particular events in the computational model, and continues to exist until it is deleted in response to other events.
- It is created in a computational context known as a Class extent, and remains in that extent for its lifetime; see Section 4.6, “Extents,” on page 4-9.
- It can have component values defined using M2-level Attributes; see Section 4.4, “Semantics of Attributes,” on page 4-4.
- It can be linked to other Instances; Section 4.7, “Semantics of Associations,” on page 4-12.

Not all M2-level Classes can have corresponding M1-level Instances. In particular, Instances can never be created for Classes which have “isAbstract” set to true. In addition, if an M2-level Class has “isSingleton” set to true, only one Instance of the class can exist within an extent for the Class.

4.4 Semantics of Attributes

Attributes are one of two mechanisms provided by the MOF Model for defining relationships between values at the M1-level. An Attribute of an M2-level Class defines a relation between each M1-level Instance of the Class, and values of some other type. The Attribute specification consists of the following properties:

- the Attribute’s “name”,
- the Attribute’s “type” which may be expressed using a Class or DataType,
- a “scope” specification,
- a “visibility” specification,
- a “multiplicity” specification,
- an “isChangeable” flag,
- an “isDerived” flag, and
- an (implicit) aggregation specification.

Many aspects of the M1-level computational semantics of Attributes depend on the mapping used. The following subsections describe those aspects of the semantics that are mapping independent.

4.4.1 Attribute name and type

The “name” and “type” of an Attribute define the basic signature of a notional binary relationship between a Class instance and an attribute value or values. For example, an Attribute declaration of the form:

```
class Class1 {
    attribute attr_1 AttrType;
};
```

defines a notional relation between an M1-level type corresponding to the Class1 and an M1-level type corresponding to the AttrType. The three main kinds of relation that can exist between for a Class and an Attribute are illustrated below in Figure 4-1. The figure shows cases where an Attribute’s “multiplicity” bounds are “[1..1]” (single-valued), “[0..1]” (optional) and “[m..n]” (multi-valued) respectively. Each notional relation is distinguishable from others for that Class by the Attribute’s “name.”

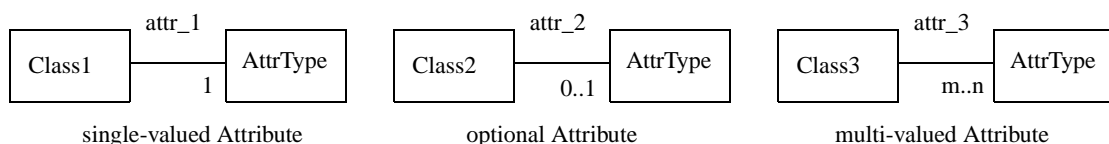


Figure 4-1 Notional Class — Attribute Relations

An M2-level Attribute's "type" can be either a Class or a DataType. In the former case, the Class — AttrType relation relates M1-level Instances corresponding to the two Classes. In the latter case, it relates M1-level Instances corresponding to the Class to M1-level Instances corresponding to the DataType.

In the following sections, it is often necessary to talk about the type of the M1-level Instances on the AttrType end of a Class — AttrType relation. To make the text more readable, we will use the phrase "the Attribute's M1-level *base type*" for this type rather than referring to it as "the M1-level type corresponding to the M2-level Attribute's "type". As we shall see, the phrase "the Attribute's M1-level *type*" is best used for another purpose.

4.4.2 Multiplicity

The "multiplicity" property defines the cardinality, uniqueness and orderedness of an Attribute as follows:

- The "lower" and "upper" fields set the bounds on the number of elements (i.e. cardinality) allowed in an Attribute value; i.e. the "(collection of) AttrType" in Figure 4-1 and Figure 4-2 above. Discussion of multiplicity usually need to deal with three cases:
 - If the "lower" and "upper" are both 1, the Attribute is single-valued; i.e. the "value" is a single instance belonging to the Attribute's M1-level base type.
 - If the "lower" is 0 and "upper" is 1, the Attribute is optional; i.e. the "value" is either an instance belonging to the Attribute's M1-level base type, or nothing.
 - Otherwise, the Attribute is multi-valued; i.e. its "value" is a collection of instances belonging to the Attribute's M1-level base type.
- The "isUnique" flag specifies whether or not a multi-valued Attribute is allowed to contain duplicates; i.e. elements that are equal according to the definition in Section 4.2.1, "Semantics of Equality for MOF Values," on page 4-3.
- The "isOrdered" flag specifies whether or not the order of the elements in a multi-valued Attribute are significant.

The "multiplicity" settings of an M2-level Attribute have considerable influence on the M1-level Attributes values. In particular, it determines whether the M1-level type of the Attribute is the M1-level base type, or a collection of that type. In addition, the "multiplicity" may also cause:

- runtime checks to ensure that a multi-valued Attribute's cardinality lies within a given range,
- runtime checks to ensure that a multi-valued Attribute does not contain duplicate members, and
- mechanisms which allow the user to specify the order of the elements of a multi-valued Attribute.

The "multiplicity" may also have considerable impact on the APIs that a mapping provides for accessing and updating Attribute values.

It should be noted that when an M2-level Attribute has “isOrdered” set to true, the corresponding Class — AttrType relation has an associated partial ordering when viewed from the Class role.

4.4.3 Scope

The “scope” of an Attribute can be either “instance_level” or “classifier_level”. For an “instance_level” Attribute, independent relationships exist between instances of MyClass and instances of AttrType. For a “classifier_level” Attribute, a single instance of AttrType (or a collection of AttrType) is related to all instances of MyClass in the extent. This is illustrated in Figure 4-2.

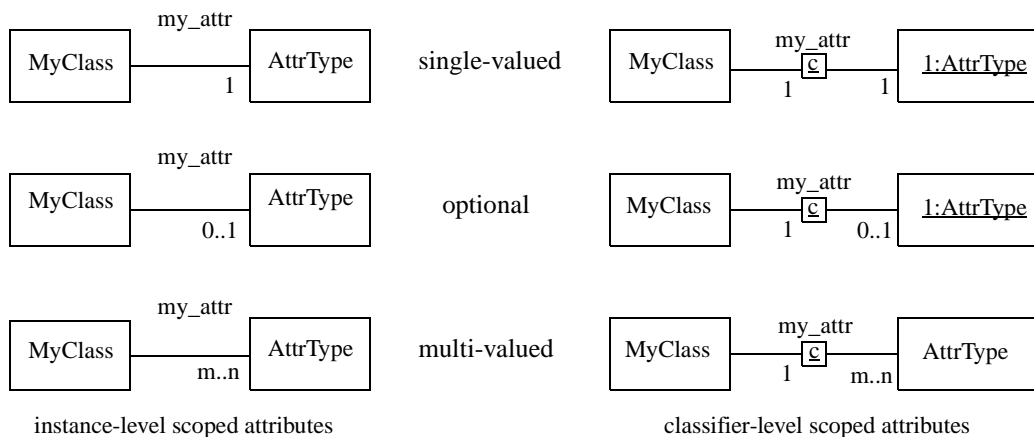


Figure 4-2 Instance-level versus Classifier-level scoping

Note – For the classifier-level Attributes, the diagrams are intended to show that all MyClass instances are related to a single instance or collection of instances of AttrType.

4.4.4 Is_derived

The “isDerived” flag indicates whether the notional relationship between a Class instance and the Attribute type instances is stored or computed.

4.4.5 Aggregation

The possible aggregation semantics of an Attribute depend on its type:

- If an Attribute’s type is expressed as a DataType, it has “non-aggregate” semantics.
- If an Attribute’s type is expressed as a Class, it has “composite” semantics.

In cases where an Attribute has “composite” semantics, the Class instance which is the value of the Attribute is a component of the Class instance which contains the Attribute, not vice-versa.

Note – The above description reflects the fact that the Attribute model element does not have an “aggregation” attribute. A Class-valued Attribute with “non-aggregate” semantics is currently expressed by making the Attribute’s type a DataType, where the DataType’s “typeCode” is an object reference type that is linked to the Class via a TypeAlias.

4.4.6 *Visibility and is_changeable*

The “visibility” property of an Attribute determines whether or not any operations for the notional relation should be present. Similarly, the “isChangeable” property determines whether update operations are present. The presence or absence of these operations do not alter the semantics of the Attribute.

4.5 *Package Composition*

This section summarizes the meta-model composition mechanisms supported by the MOF Model and discusses their impact on M1-level semantics.

4.5.1 *Package Nesting*

Package nesting is the simplest of the MOF’s Package composition mechanisms. At the M2-level, Package nesting is expressed by making the outer Package the “container” of the nested Package. The definition of the Contains association in the MOF Model means that Package nesting is a strict composition relationship.

The main intended function of Package nesting is information hiding. Placing a Class or DataType in an inner Package rather than an outer one notionally makes it less visible to other meta-models. When the MOF visibility rules are defined (in a future revision of this specification), this information hiding will be more strongly enforced.

Nesting of Packages also affects the M1-level interfaces and implementations. The meaning of any element of a meta-model is potentially dependent on its context in a variety of ways. Thus, when the element is defined in a nested Package, its meaning may depend on the outer Package; e.g. on Constraints or Classifiers declared therein. This means that anything that uses a nested element will also implicitly depend on the context. To avoid potential M1-level anomalies caused by this kind of dependency, the MOF Model does not allow a meta-model to import a nested Package or a Classifier defined within a nested Package.

The M1-level semantics of Package nesting are as follows. The behaviour of an M1-level instance of a Classifier declared in a nested Package depends on state in both its immediate Package, and its enclosing Packages. As a result, the M1-level instance of the nested Classifier is inextricably tied to other instances within the outermost enclosing Package extent; see Section 4.6.4, “Package Extents,” on page 4-11.

4.5.2 *Package Generalization*

Package generalization allows an M2-level Package to reuse all of the definitions of another M2-level Package. Package generalization is expressed at the M2-level by connecting the super-Package and sub-Package using a Generalizes link. (The MOF Model's Constraints mean that Generalization is effectively an aggregation in the UML sense.)

The M1-level semantics of Package generalization are as follows. The behaviour of M1-level instances of the elements of an M2-level Package typically depends on M1-level behaviour for M2-level super-Package elements. Indeed, an M1-level Package "instance" is substitutable for M1-level Package instances for M2-level super-Packages.

Package inheritance does not create any relationship between an instance of the super-Package and an instance of the sub-Package. Therefore an M1-level Package extent is not related to M1-level super- or sub-Package extents; see Section 4.6.4, "Package Extents," on page 4-11.

4.5.3 *Package Importation*

Package importing allows an M2-level Package to selectively reuse definitions from other M2-level Packages. Package importation is expressed at the M2-level by placing an Import in the importing Package that is related to the imported Package by an Aliases link. In this case, the M2-level Import object has its "isClustered" attribute set to false. Since Package importation can be cyclic, it is neither an aggregation or a composition in the UML sense.

Note – The MOF Model's Constraints make it illegal for a Package to import itself, or for any Package to import a nested Package. Furthermore, while the MOF Model allows Package importation to be cyclic, the preconditions for the MOF Model to IDL mapping disallow most dependency cycles, including those between Packages that result from cyclic importation.

The M1-level semantics of Package importation are minimal. No substitutability or state relationships exist between the M1-level instances of an importing or imported Package, or between their respective extents. Indeed, an importing Package will typically not even share implementation code with the imported Package.

4.5.4 *Package Clustering*

Package clustering allows an M2-level Package to selectively reuse definitions from other M2-level Packages, and also share M1-level implementation infrastructure. The M2-level expression of Package clustering is similar to that for Package importation; see above. The difference is that the Import object has "isClustered" set to true.

The M1-level semantics of Package clustering are similar to those of Package nesting because a cluster Package instance has its clustered Package instances as its components. However, unlike nested Packages, it is still possible to have a free-standing M1-level instance of such a Package whose extent is unrelated to any extent of a cluster Package.

4.6 *Extents*

This section introduces the concept of an “extent” in more detail, and then gives the formal definitions of the extent of a Class, an Association and a Package.

4.6.1 *The Purpose of Extents*

Current generation middleware systems typically aim to allow clients to use objects without knowledge of their locations or context. However, groups of objects generally exist in the context of a “server” which has responsibility for managing them. The implementation an object often uses knowledge of its shared context with other objects to optimize performance, and so forth.

While statements about object location have no place in the MOF specification, the MOF Computational Model assumes a notion of context in many areas:

- The classifier-scoped features of an M2-level Class are notionally common to “all instances” of the Class.
- Mappings typically allow a client to query over “all links” in an Association instance.

It is impractical to define “all instances” or “all links” as meaning all instances or links in the known universe. Therefore, MOF specification defines logical domains of M1-level instances that are the basis of these and other “for all” quantifications. These domains of M1-level instances are called extents.

Figure 4-3 below shows the extents defined by two “instances” (on the right) of the example meta-model on the left. Notice that the static nesting of Packages, Classes and Associations inside other Packages is mirrored in the extents (i.e., the dotted ovals).

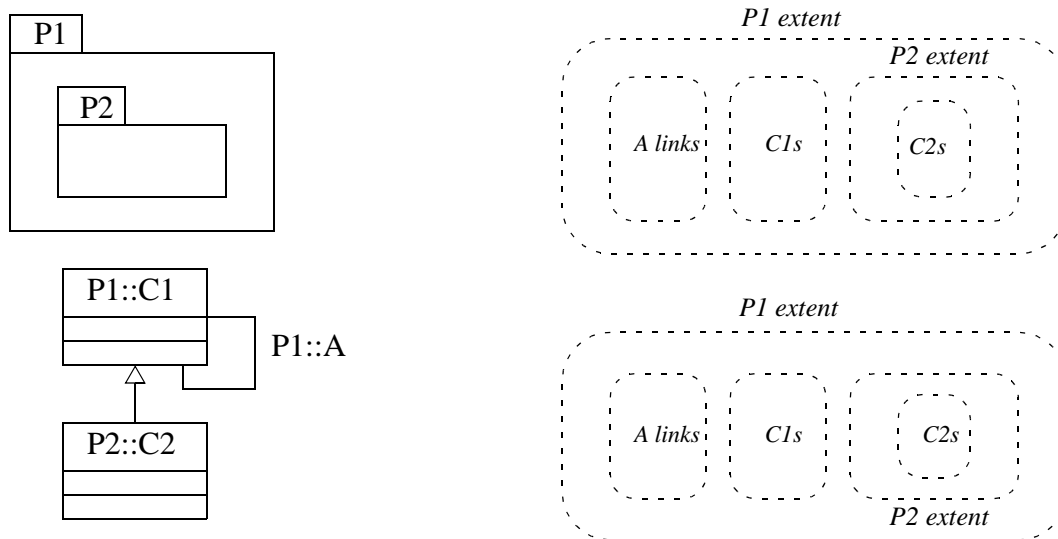


Figure 4-3 Extents for two meta-model instances

Every Class instance or link belongs to precisely one Class or Association extent. These extents are part of Package extents, depending on the “lexical” structure of the meta-model. This means that extents are strictly hierarchical in nature. As we shall see in Section 5.2.1, “Meta Object Type Overview,” on page 5-2, extents are related to the intrinsic container semantics of meta-objects.

Note – There is no requirement that extents have any manifestation in the partitioning of objects between physical MOF servers. However, there are clear performance advantages in implementing such a partitioning.

4.6.2 Class Extents

The extent of a Class is defined to be the complete set of M1-level instances of the Class that share classifier-scoped properties (e.g., Attribute values). A Class instance is created in the context of a Class extent and remains within that extent for its entire lifetime (i.e., until the instance is explicitly deleted).

4.6.3 Association Extents

The extent of an Association is defined to be the complete set of M1-level links for the Association. A link is created in the context of an Association extent and remains within that extent for its entire lifetime.

4.6.4 Package Extents

The extent of a Package is a conglomeration of the extents of Classes, Associations and other Packages according to the following rules:

1. When an M2-level Package contains a Class or Association, an extent for the Package contains extents for the Classes and Associations.
2. When an M2-level Package contains nested Packages, an extent for the outer Package contains extents for the inner Packages.
3. When an M2-level Package clusters one or more other Packages, an extent for the cluster Package aggregates the extents for the clustered Packages.
4. When an M2-level Package inherits from another Package, an extent for the sub-Package:
 - a. contains an extent for each nested Package, Class or Association in the super-Package
 - b. aggregates an extent for each Package clustered by the super-Package, and
 - c. aggregates or contains extents by recursive application of rule Chapter 4. to the super-Package's super-Packages.

When a Package inherits from another Package by more than one route, the sub-Package extent will contain one extent for each directly or indirectly inherited Class, Association or nested Package. This is illustrated in Figure 4-4 below. Notice that the extent for Package P4 contains only one C1 extent.

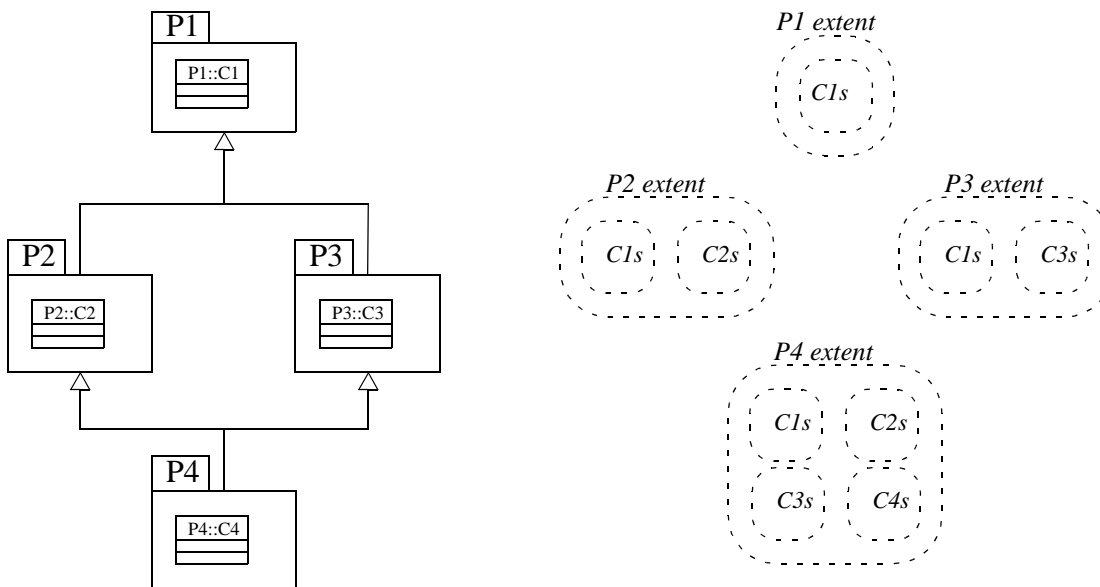


Figure 4-4 Extents for Multiply Inheriting Packages

When a Package clusters other Packages by more than one route, the outer cluster Package will contain one extent for all directly or indirectly clustered Packages. This is illustrated in Figure 4-5 on page 4-13. Notice that the relationship between the extents of a cluster Package and the extents of the clustered Packages is aggregation rather than strict containment. In particular, in the P4 case, the extent for P1 is not fully contained by either the P2 or P3 extents.

Note – The extent for an M2 Package that imports (rather than clusters) other Packages does *not* contain extents for the imported Packages or their contents.

4.7 Semantics of Associations

Associations are the MOF Model’s second mechanism for relating MOF values at the M1-level. A MOF M2-level Association defines a binary relation between pairs of M1-level Instances, where the relationships in the relation are called Links. The Links for a given M2-level Association conceptually belong to a Link set.

Note – While the MOF Model appears to support N-ary Associations, this is not so. There is a Constraint that states that an Association must have precisely 2 Association Ends; see “AssociationsMustBeBinary” on page 3-104.

An M2-level Association definition specifies the following properties:

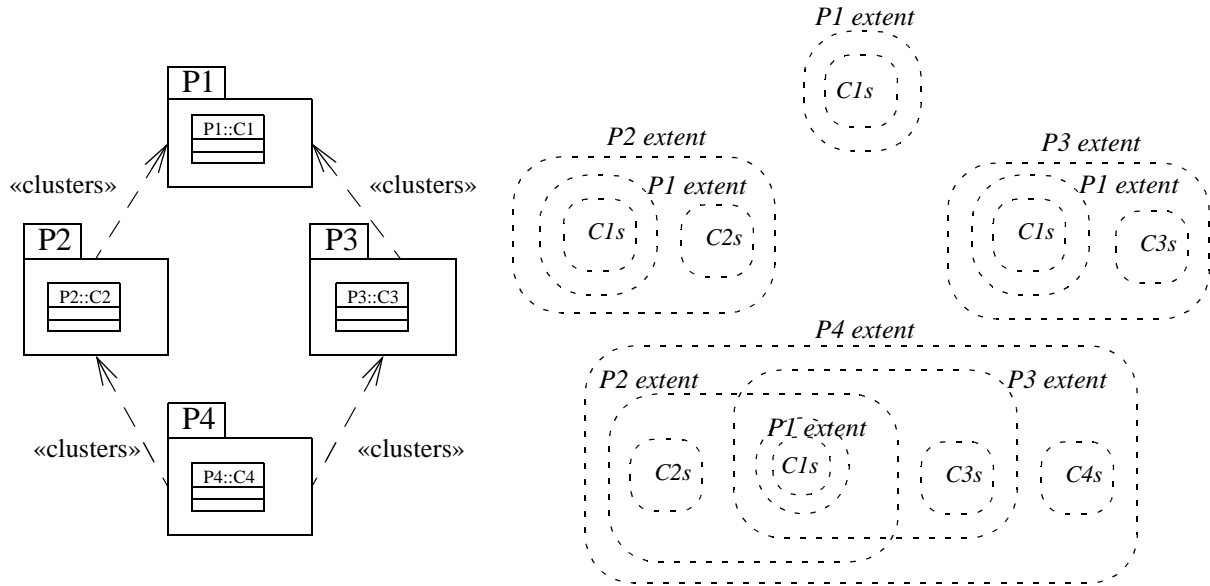


Figure 4-5 Extents for Clusters of Clusters

- an Association “name”,
- a pair of AssociationEnds which each have:
 - a “name”,
 - a “type” which must be a Class,
 - a “multiplicity” specification,
 - an “isNavigable” flag, and
 - an “isChangeable” flag.
- an “isDerived” flag which determines whether the Association Links are stored explicitly or derived from other state.

4.7.1 MOF Associations in UML notation

A MOF Association is represented in UML notation as shown in Figure 4-6 below.

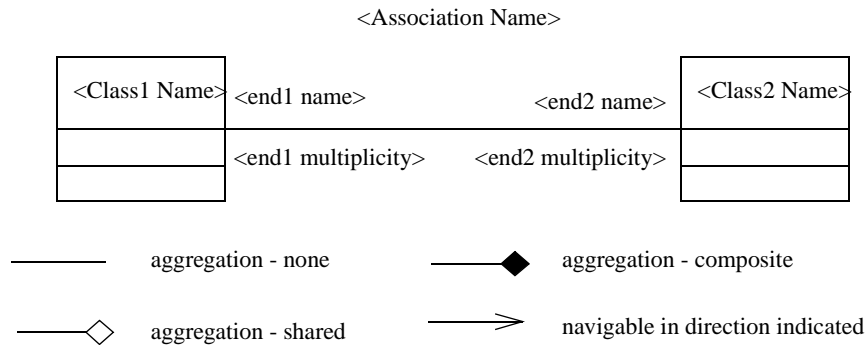


Figure 4-6 An M2-level Association in UML notation

The connecting line denotes an Association between two Classes. The text of <Association Name>, <end1 name> and <end2 name> denote the “name” values for the respective Association and AssociationEnds. If the Association name is preceded by a forward slash, the Association has “isDerived” set to true.

The Class boxes denote the respective types of the two ends. If the two ends of an Association have the same type, the Association line loops around so that it connects a Class box to itself.

The <end1 multiplicity> and <end2 multiplicity> text give the multiplicity settings for the respective ends of the Association. The text that can appear here consists of an optional bounds specification with syntax:

```
<bounds> ::= [ <number> \..\' ] ( <number> | \*' )
```

and the optional keyword “ordered.”

Finally, the navigability and aggregation of the ends of the Association are (partially) specified by the symbols at the respective ends of the line:

- An empty diamond indicates that the Instances at the labeled end “shares” the Instances at other end.
- A filled diamond indicates that the Instances at the labelled end are “composed” of Instances at the other end.
- An arrow head indicates that the Association is navigable from the Instance at the other end to the Instance at the labelled end.

Note – There are a couple of anomalies in the mapping of UML Association notation to MOF Associations. First, while navigability and aggregation are orthogonal in the MOF, it is not possible to put both a diamond and an arrow head on the same end of a UML Association line. This means, for example, that it is not possible to express (the lack of) navigability from a component end to a composite end. Second, UML is imprecise about what an Association line with no arrowheads means. It can mean that the Association is not navigable, or alternatively that its navigability is not shown.

4.7.2 Core Association Semantics

This section defines the core semantic model for M1-level Association instances in a rigorous, mapping independent fashion, and enumerates some important characteristics that follow from the definition.

4.7.2.1 A Mathematical Model of Association State

Given an M2 Association labelled as in Figure 4-6 above, the mapping to M1-level Link sets and Links can be modeled as follows:

1. The M1-level Instances of the M2-level Classes <Class1> and <Class2> belong to sets *Class1_Instances* and *Class2_Instances* that represent the sets of all possible instances of <Class1> and <Class2>. (Note these sets are not restricted to current extant instances.)
2. The set *All_Links* is the Cartesian product of the sets *Class1_Instances* and *Class2_Instances*. Thus a *Link*, which is a member of *All_Links*, can be any tuple of the form “<c1, c2>” where “c1” and “c2” are members of *Class1_Instances* and *Class2_Instances* respectively.
3. The *Link_Set* is a subset of the set *All_Links* which consists of those *Links* that currently exist in the given M1-level Association.
4. If one or other of the AssociationEnds has “isOrdered” set to true, there is a partial ordering *Before* over the elements of *Link_Set* defined as follows. Assuming that <End1> of the Association is the one that is flagged as ordered:
 - a. For each Instance “i” in *Class2_Instances*, we can define a subset *End2_Links_i* of *Link_Set* consisting of those *Links* in *Link_Set* for which the second tuple member is “i”.
 - b. Given the *End2_Links_i* sets as defined in item Chapter a. above, the *Before* ordering is defined between any pair of different *Links* in a *End2_Links_i* set with 2 or more members. In other words, for any distinct *Link_j* and *Link_k* in *End2_Links_i*, we can say either *Link_j Before Link_k*, or *Link_k Before Link_j*.
 - c. The *Before* ordering is NOT defined between any pair of *Links* that belong to different *End2_Links* sets.
 - d. Where it is defined, the *Before* ordering is required to be:
 - i. transitive; i.e. *L_i Before L_j* and *L_j Before L_k* implies that *L_i Before L_k*, and

- ii. anti-reflexive; i.e. $L_i \text{ Before } L_j$ implies $\text{not } L_j \text{ Before } L_i$.

(If $\langle \text{End2} \rangle$ of the Association is ordered, substitute End2 for End1 and vice versa in the above.)

5. A *State* of an M1-level instance of an Association consists of the *Link_Set* and (if the Association is ordered) the *Before* ordering
6. A *Well-formed State* is a *State* in which:
 - a. The *Links* set is a subset of *Valid_Links*, where *Valid_Links* is the subset of *All_Links* where the connected Instances currently exist.
 - b. The *End_Links_i* sets as defined in item Chapter 4.Chapter a. above conform to their respective Association End upper and lower bounds; i.e.
 - i. the number of *Links* in each *End1_Links_i* set must be greater than or equal to $\langle \text{End2.lower} \rangle$, and less than or equal to $\langle \text{End2.upper} \rangle$, and
 - ii. the number of *Links* in each *End2_Links_i* set must be greater than or equal to $\langle \text{End1.lower} \rangle$, and less than or equal to $\langle \text{End1.upper} \rangle$.

Ideally, the computational semantics of M1-level Associations for a particular mapping should be describable as transformations from one *Well-formed State* to another. However, some mappings must be defined such that the *State* of an Association instances is not always a well-formed. For example, in the IDL mapping, deletion of an Instance may cause a *End_Links* set to contain too few *Links*.

The general model of an M1-level Association's *State* may be further constrained by M2-level Constraints on the Association or other elements of the meta-model. Other systematic restrictions may apply in some mappings; e.g. "The Reference Closure Rule" on page 4-19 and "The Composition Closure Rule" on page 4-20.

4.7.2.2 Characteristics of M1-level Associations.

The definitions of *Links* and *Link_Sets* above mean that M1-level Association instances have the following characteristics:

- *Links* only exist between existing Instances in a *Well-formed State*. When an Instance ceases to exist, any *Links* involving the Instance in any *Link_Set* cease to be universally meaningful.
- A *Link* " $\langle a, b \rangle$ " is distinct from a *Link* " $\langle b, a \rangle$ ". In other words, *Links* are directed. (Whether or not the "direction" of a *Link* has a meaning depends on the underlying semantics of the reality that the M2-level Association describes.)
- *Links* do not have object identity, but are uniquely identified by the Instances at both ends.
- Since a *Link_Set* is defined to be a set, it cannot contain more than one copy of a given *Link*. In other words, M1-level Associations cannot contain duplicate links.
- The *Before* ordering on the *Links* in an *End_Links* set (where defined) can be represented by arranging the *Links* in a strictly linear sequence.

- There can be multiple *States* for a given M2-level Association, each corresponding to a different M1-level Association instance in separate Package instances. In this scenario:
 - a given *Link* can be a member of multiple *Link_Sets*, and
 - the *Before* orderings of different *States* will be independent.

4.7.3 AssociationEnd Changeability

The “isChangeable” flag for an AssociationEnd determines whether or not the APIs for the Association should allow clients to change Links in an M1-level Association instance. The precise interpretation of this flag is mapping specific.

4.7.4 AssociationEnd Navigability

The “isNavigable” flag for an AssociationEnd determines whether or not client should be able to “navigate” the Links in an M1-level Association instance. The flag also determines whether or not the AssociationEnd can be used as a “key”. This flag’s interpretation (i.e. its impact on APIs) will depend on the mapping used.

4.7.5 Association Aggregation

The “aggregation” attributes of an Association’s two ends determines the aggregation semantics for the corresponding M1-level Association instances; see Section 4.8, “Aggregation Semantics,” on page 4-17. The impact of aggregation semantics are largely mapping specific. However, “composite” aggregation does place constraints on the *Link_Set* of a *Well-formed State*.

4.7.6 Derived Associations

When an M2-level Association has “isDerived” set to true, the resulting M1-level Association’s *Link_Set* is calculated from other information in the M1-level model. The M1-level semantics of derived Association instances is beyond the scope of the MOF specification.

4.8 Aggregation Semantics

As noted previously, the MOF Model provides two ways of relating MOF values; i.e. Associations and Attributes. In both cases, a relation has a property known as aggregation that determines how strongly related values are tied together.

The MOF Model currently supports three aggregation semantics; i.e. “none”, “shared” and “composite”, in order of increasing strength.

Note – In practice, the semantics of aggregation are mostly concerned with the life-cycles of related values. Since different mappings will use different strategies for managing the life-cycles of values, aggregation semantics are largely mapping specific.

4.8.1 Aggregation “none”

An Attribute or Association with aggregation of “none” has the weakest form of relation between values. This will typically correspond to independent life-cycles for both parties and the use of shallow copy semantics in a mapping.

4.8.2 Aggregation “composite”

An Attribute or Association with aggregation of “composite” has the strongest form of relation between values. A “composite” relation involving two types is asymmetric, with one “end” labelled as the “composition” type and the other end labelled the “component” type. An instance of the first type is “composed of” instances of the second type.

An M1-level “composite” relation is defined to have information model semantics that can be loosely described as containment semantics:

1. If a value “v1” is a component of some other value “v2” in a given composite relation, “v1” may not be a component of any other value “v3” in any composite relation. In short, a value can have *at most* one container in any “composite” relation.
2. A value may not be related to itself in the closure of any “composite” relations. In short, a value may not directly or indirectly contain itself.

Other restrictions may apply to “composite” relations in some mappings; e.g. “The Composition Closure Rule” on page 4-20.

4.8.3 Aggregation “shared”

An Attribute or Association with aggregation of “shared” corresponds to a relation between values that is between “none” and “shared”.

Note – The semantics of “shared” aggregation should correspond to the semantics of an Aggregate in UML. Unfortunately, the OMG UML specification gives no clear guidance on what these semantics should be. As an interim measure, the use of “shared” aggregation in the MOF is discouraged.

4.9 Closure Rules

The MOF’s support for multiple Package “instances” introduces some potential anomalies into the computational model. These are resolved by three “closure” rules based on the definitions of extents in Section 4.6, “Extents,” on page 4-9.

4.9.1 The Reference Closure Rule

Recall that a Reference value is defined as a projection of an M1-level Class instance in an Association. Given that Association link sets are not global, a reference's value must be a projection in a particular link set. There is an "obvious" candidate link set for typical M1-level Class instances, namely the link set belonging to the Package "instance" that contains the Class instance. This is shown in Figure 4-7.

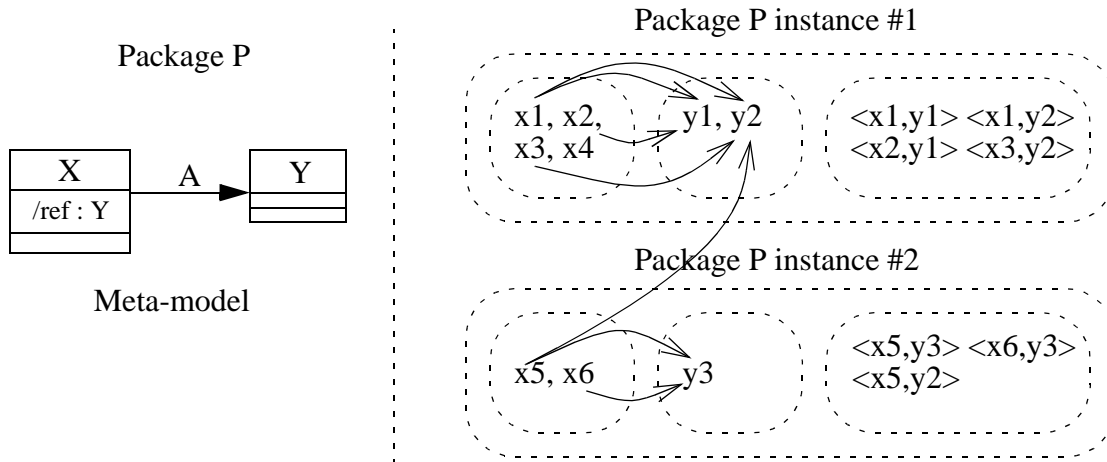


Figure 4-7 References for multiple Package instances

Figure 4-7 shows the Y instances visible to each X instance in two Package instances. Notice that the link set in the second Package instance contains a link to a Y instance belonging to the first Package instance (i.e., "**<x5,y2>**"). This presents no particular problems, since the "x5" object can find the link to "y2" by looking in the A link set for its containing Package instance.

However, suppose that the "**<x5,y2>**" had been in the A link set for the first Package instance. Now an instance of the X Class has to look in the link sets of both (or in the general case, all) Package instances to find all of the links. Alternatively, an X instance might only look in the link set for its owning Package instance, leading to non-intuitive computational semantics for Reference values. (Consider the case where there are References for both Association Ends.)

To avoid such non-intuitive (and arguably anomalous) semantics, the computational semantics for Associations includes a runtime restriction that prevents the problematic links from being created. This restriction is called the *Reference Closure Rule*:

"If Class C has a Reference R that exposes an Association End E in an Association A, then it is illegal to cause a link to be constructed such that an instance of C (or a sub-class of C) at the exposed End belongs to a different outermost extent to the A link set containing the link".

The *Reference Closure Rule* is shown graphically by Figure 4-8 below for the case of an Association with a Reference to one end. The Reference Closure Rule is enforced by runtime checks on M1-level operations that construct links (e.g., the link add and

modify operations). This can be achieved by using the “outermost_containing_package” operations on the respective meta-objects; see Section 6.2, “The Reflective Interfaces,” on page 6-3.

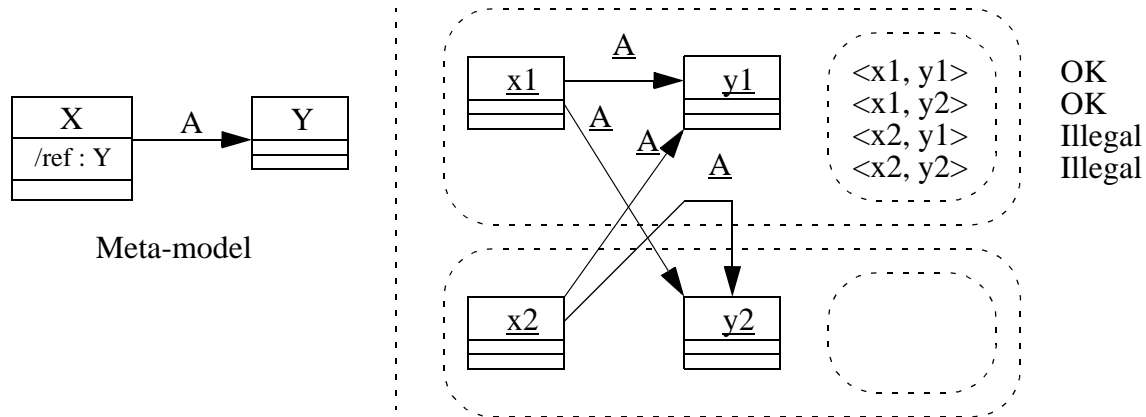


Figure 4-8 The Reference Closure Rule

4.9.2 The Composition Closure Rule

The MOF Model provides constructs for declaring that the instances of one meta-model element are “composed of” instances of another; see Section 4.8, “Aggregation Semantics,” on page 4-17.

One of the key properties of composites is that a composite instance and its component instances have the same lifetime; i.e. when a composite meta-object is deleted, all of its components are also deleted. This is not difficult to implement when the composite instance and its components all belong to the same Package instance. However, a range of problems can arise when a composition crosses one or more outermost Package extent boundaries. For instance:

- How do the server implementations for the respective extents ensure that deletion is reliable in the face of server crash, network partition and so on?
- What are the access control implications of compositions? For example, should a client of one server / extent be able to implicitly delete components held in another server / extent?

To avoid having to deal with these difficult questions, the MOF computational model restricts the situations in which compositions may be formed. This restriction is called the *Composition Closure Rule*:

“The composite and component instances in a composition along with any links that form the composition must all belong to the same outermost Package extent.”

The *Composition Closure Rule* is shown graphically by Figure 4-3. This shows the rule as it applies to both composite Attributes and composite Associations.

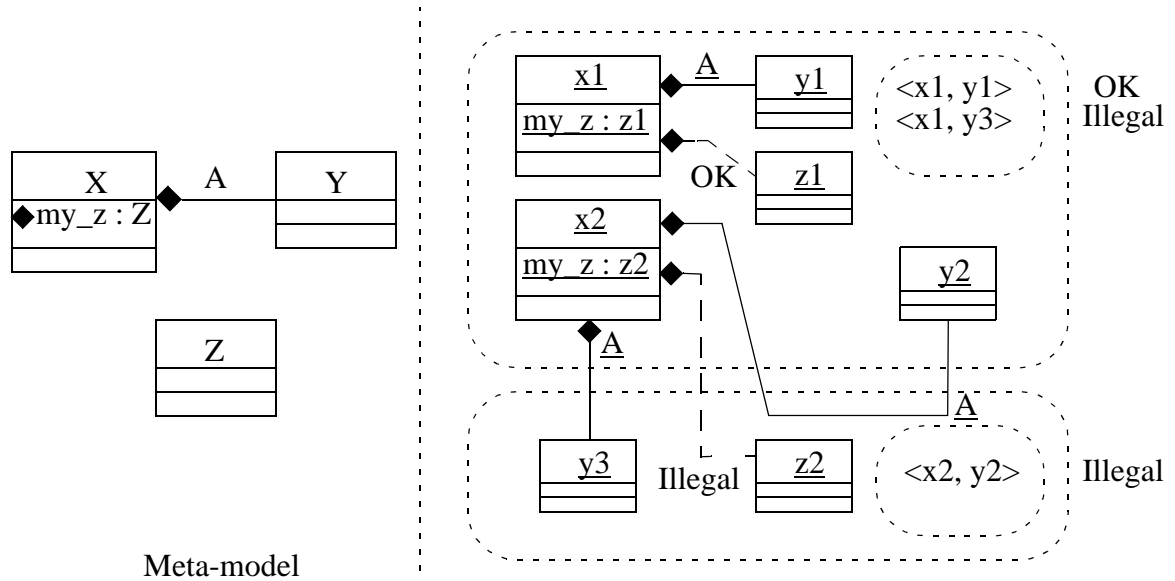


Figure 4-9 The Composition Closure Rule

The Composite Closure Rule is enforced by runtime checks on M1-level operations that construct links in an Association with Composite semantics; e.g. the link add and modify operations. Similar checks are required for operations that update composite Attributes. The checks can be implemented by using the “immediate_copntainer” and “outermost_containing_package” operations on the relevant meta-objects; see Section 6.2, “The Reflective Interfaces,” on page 6-3.

4.10 Recommended Copy Semantics

It is envisaged that some MOF mappings will provide APIs for copying metadata. The purpose of this section is to recommend a semantic model for such copy operations. Suggested semantics are given for “shallow” and “deep” copying. (A shallow copy is one in which conceptual components of an object are copied and other connected objects are not. A deep copy is one in which both components and more loosely related objects are copied.)

The following table details what objects should and should not be copied. The semantics are defined from the perspective of an object being copied.

Table 4-1 Copy semantics for different kinds of relationships.

Construct	Target type	Aggregation	Shallow Copy	Deep Copy
Attribute	Instance	none	The Attribute value in the copy will be the same Instance value as in the original.	The Attribute value in the copy will be the same Instance value as in the original.
Attribute	MOF data type	none	The Attribute value in the copy will be the same data value as in the original. Embedded Instance values will be the same as in the original.	The Attribute value in the copy will be the same data value as in the original. Embedded Instance values will be the same as in the original.
Attribute	Instance	composite	The Attribute value in the copy will be a shallow copy of the Instance value as in the original.	The Attribute value in the copy will be a deep copy of the Instance value in the original.
Association	Instance	none	No link is created.	A link is created from the copy to the original link target.
Association	Instance	shared	A link is created from the copy to the original link target.	A link is created from the copy to a deep copy of the original link target.
Association	Instance	composite	A link is created from the copy to a shallow copy of the original link target.	A link is created from the copy to a deep copy of the original link target.

Unless otherwise stated, copying of a group of Instances related by Association or Attributes should give a 1-to-1 mapping between original Instances and copied Instances, and their respective relationships.

Note – The above suggested semantics do not cover copying of MOF values whose type is a native type. Those semantics will depend on whether or not the values in question are copyable.

4.11 Computational Semantics

4.11.1 A Style Guide for Metadata Computational Semantics

While the MOF specification gives the required computational semantics for M1-level metadata, it does not (and should not) state that these semantics constitute the only behavior. It is envisaged that vendor and end-user implementations of metadata servers may support additional semantics. In addition, the computational semantics of M2-level derived Attributes, derived Associations and Operations are not specified at all in the standardized part of the MOF Model.

In theory, the complete computational semantics of a meta-model server can include any behavior that the implementer chooses. The purpose of the section is to set down some conventions to guide the implementer.

4.11.2 *Access operations should not change metadata*

Many operations on Instance and Associations are provided to support access to the public state of a model; e.g. the “get” operations for Attributes, the “query” operations for Associations. For normal (non-derived) Attributes and Associations, the standard computational semantics of an access operations are to simply return the corresponding value or collection. For derived Attributes and Associations, there are no standard semantics at all.

In general, it is bad style for an access operation to have observable side-effects on the primary metadata. Similarly, it is bad style for an Operation with “isQuery” true to have such side-effects.

The rationale for this rule is that the user would not expect an access operation to have visible side-effects.

Note – It may be reasonable (for example) for an Attribute “get” operation to update a private counter Attribute that records the number of accesses. The legitimacy of this kind of behavior depends on whether or not the state modified can be classified as “primary” metadata.

4.11.3 *Update operations should only change the nominated metadata*

The standard semantics of metadata update operations define which metadata is expected to be modified by the operation. However, there is no explicit requirement that other metadata should not be changed.

It is bad style for an update operation for a non-derived Attribute, Reference or Association to change any primary metadata other than that which is identified by the standard semantics.

The rationale for this rule is that the user would not expect such changes to occur.

Note – This rule is not intended to apply to operations for derived Attributes, References or Associations, or to Operations with “isQuery” false.

4.11.4 *Derived Elements should behave like non-derived Elements*

M2-level Attributes and Associations can be defined as being derived from other information in a meta-model (i.e., by setting the respective “isDerived” flag to true). The required M1-level behaviour of derived Elements is identical to that for equivalent non-derived Elements. Behavior that contradicts the semantics in this chapter and in the relevant mapping specification is non-conformant.

However, since derived Attributes and Associations have to be implemented using mechanisms that are beyond the scope of the MOF Model, conformance is ultimately the responsibility of the meta-model implementer.

It is recommended that implementer defined M1-level operations for derived Elements should have MOF conformant behavior. The alternative (non-conformant behavior) tends to break the illusion that the Attribute or Association is “real,” and should be avoided. If the required semantics are unimplementable, the meta-model is incorrect.

4.11.5 Constraint evaluation should not have side-effects

The MOF specification does not define how Constraints defined in a meta-model should be evaluated. In particular, it does not define whether Constraint evaluation can change the metadata.

It is bad style for the evaluation of a Constraint to change metadata.

The rationale is two fold. First, Constraints are provided as mechanism for specifying metadata correctness, not as a mechanism for defining behavior. Second, since the MOF specification does not say when Constraint evaluation should occur (in all cases), side-effects in Constraint evaluation could be a major source of interoperability problems.

4.11.6 Access operations should avoid raising Constraint exceptions

The MOF specification does not define when deferred Constraint evaluation should occur. In theory, it can occur at any time, including when the user invokes an access operation.

It is bad style for an access operation on a non-derived Attribute, Reference or Association to raise an exception to indicate that the metadata is structurally inconsistent or that a Constraint has been violated.

The rationale is that an application program that is reading metadata (rather than updating it) is typically not in a position to do anything about the violation of deferred structural constraints or model specific Constraint. Alternatively, an application may try to access the metadata, knowing that it is inconsistent, so that it can then correct it.

It is bad style for an access operation on a derived Attribute, Reference or Association to raise a similar exception *unless* the inconsistency makes it impossible to calculate the required derived value(s). The same rule applies to Operations with “isQuery” true.

The rationale being less prescriptive about derived access operations is that the formulae used to derive the value(s) will typically have certain assumptions about the consistency of the metadata.

Contents

This chapter contains the following topics.

Topic	Page
“Overview”	5-1
“Meta Objects and Interfaces”	5-2
“Computational Semantics for the IDL Mapping”	5-6
“Exception Framework”	5-23
“Preconditions for IDL Generation”	5-32
“Standard Tags for the IDL Mapping”	5-35
“Generated IDL Issues”	5-38
“IDL Mapping Templates”	5-46

5.1 Overview

This chapter defines the standard mapping from a model defined using the MOF Model onto CORBA IDL. The resulting interfaces are designed to allow a user to create, update and access instances of the model using CORBA client programs. While the standard IDL mapping implies detailed functional semantics for an object server for a mapped model, it does not define the implementation.

Note that while the mapping specification is defined to be easy to automate, a conformant MOF implementation is not required to support automatic IDL generation.

5.2 Meta Objects and Interfaces

This section describes the different kinds meta-objects that represent MOF-based meta-data in a CORBA environment.

5.2.1 Meta Object Type Overview

The MOF to IDL mapping and the Reflective module share a common, object-centric model of meta-data with five kinds of M1-level meta-object; i.e. “instance” objects, “class proxy” objects, “association” objects, “package” objects and “package factory” objects. The relationships between M2-level concepts and M1-level objects is illustrated by the example in Figure 5-1.

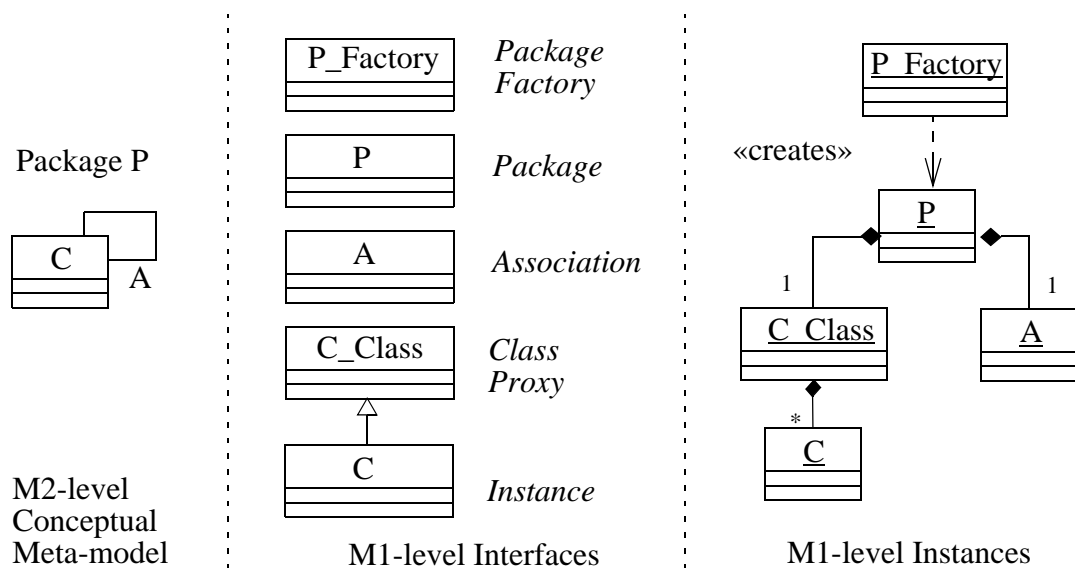


Figure 5-1 Relationships between M1 and M2 level

The example shows how a simple M2-level meta-model (on the left) maps onto the five kinds of M1-level meta-object (in the center). The right of the diagram shows the intrinsic conglomeration relationships that exist between the meta-objects in a Package “instance.” (As noted, in Section 4.6, “Extents,” on page 4-9, these relationships do not always have strict containment semantics.)

Note – These intrinsic conglomeration relationships exist for all M1-level meta-objects. They have no explicit meaning in connection with the represented meta-data. Rather, they are provided to assist the management of meta-objects. (The intrinsic conglomeration relationships should not be confused with the M1-level composition relationships that correspond to M2-level composite Associations and Attributes.)

5.2.1.1 Package objects and Package Factory objects

The instances of an M2-level Package are represented as Package objects. A Package object is little more than a “directory” of read-only attributes that give access to a collection of meta-objects described by a meta-model. The attributes of a Package object refer to “static” objects. In particular, there is

- one Package attribute for each M2-level Packages that is nested or clustered by the Package (none are present in the example above),
- one Class Proxy attribute for each M2-level Class in the Package, and
- one Association attribute for each M2-level Association in the Package.

The number and types of the static objects, and the corresponding attributes in an M1-level Package interface is determined by the M2-level Package specification. The objects cannot be directly created, destroyed, added or removed by a client.

While there is a usually a one-to-one correspondence between the Packages’ reference attributes and the static objects, this need not be the case. The correspondence is actually determined by the extent relationships as described in Section 4.6.4, “Package Extents,” on page 4-11. Thus, for example, when an M2-level Package is clustered by more than one route, there should be one M1-level Package object that is accessed via two attributes.

A Package object is typically obtained by invoking a “create” operation on a Package Factory objects. This creates the Package object, and all of the necessary static objects. The arguments to the “create” operation are used to initialize any classifier-scoped Attributes defined within the M2-level Package.

5.2.1.2 *Class Proxy objects*

As stated above, a Package object contains one (and only one) Class Proxy object for each M2-level Class in the M2-level Package. A Class Proxy object serves a number of purposes:

- it is a factory object for producing Instance objects in the Package “instance,”
- it is the intrinsic container for Instance objects, and
- it holds the state of any classifier-scoped Attributes for the M2-level Class.

The interface of a Class Proxy object provides operations for accessing and updating the classifier-scoped attribute state. Other operations allow a client to invoke classifier-scoped Operations.

The interface also provides a factory operation allows the client to create Instance objects. It also gives read-only access to the set of extant Instance objects contained by the Class Proxy object.

5.2.1.3 *Instance objects*

The instances of an M2-level Class are represented by Instance objects. An Instance object holds the state corresponding to the instance-scoped M2-level Attributes for the Class, and any other “hidden” state implied by the Class specification. Generally speaking, many Instance objects can exist within a given Package “instance”.

As described above, Instance objects are always contained by a Class Proxy object. The Class Proxy provides a factory operation for creating Instance objects that takes initial values for the instance-scoped Attributes as parameters. When an Instance object is created, it is automatically added to the Class Proxy container. An Instance is removed from the container when it is destroyed.

The interface for an Instance object inherits from the corresponding Class Proxy interface. In addition it provides:

- operations to access and update the instance-scoped Attributes,
- operations to invoke the instance-scoped Operations,
- operations to access and update Associations via Reference,
- operations that support object identity for the Instance, and
- an operation for deleting the Instance object.

5.2.1.4 *Association objects*

Links that correspond to M2-level Associations are not represented as meta-objects. Instead, an M1-level Association object holds a collection of links (i.e. the link set) corresponding to an M2-level Association. The Association object is a “static” object that is contained by a Package object, as described previously. Its interfaces provides:

- operations for querying the link set,
- operations for adding, modify and removing links from the set, and
- an operation that returns the entire link set.

5.2.2 *The Meta Object Interface Hierarchy*

This section describes the patterns of interface inheritance in the CORBA IDL generated by the MOF to IDL mapping. The patterns are illustrated in Figure 5-2 below.

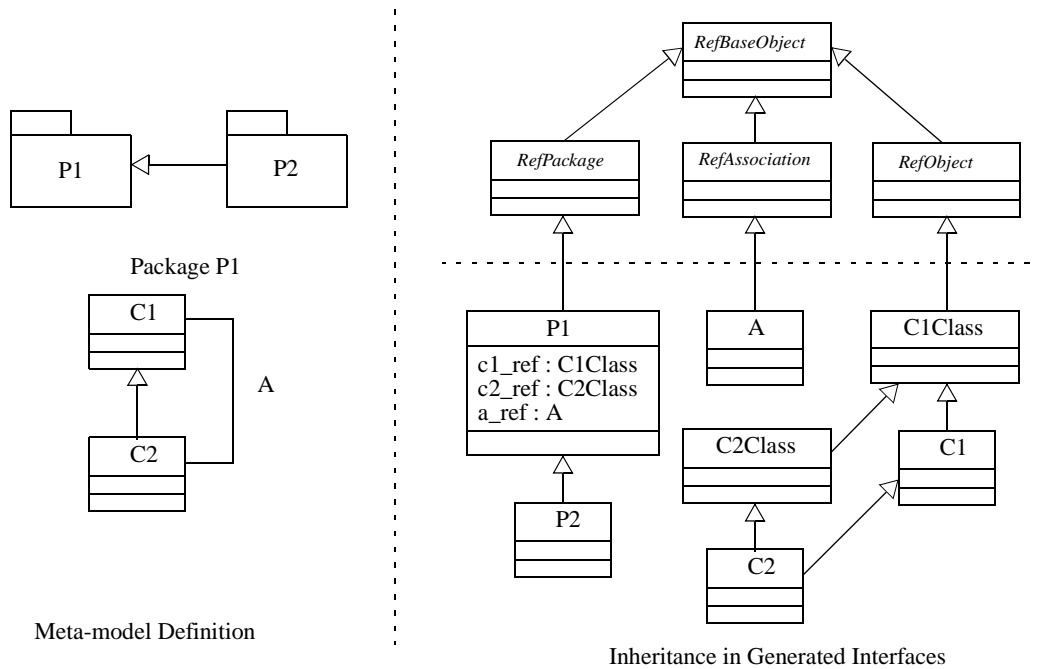


Figure 5-2 Generated IDL Inheritance Patterns

Figure 5-2 shows an example MOF meta-model expressed in UML (on the left) that consists of two Packages P1 and P2. The first Package P1 contains Classes C1 and C2, where C2 is a subclass of C1 and an Association A that connects C1 and C2. The second Package P2 is then defined as a subpackage of P1.

The UML class diagram (on the right) shows the inheritance graph for the generated interfaces corresponding to the example meta-model.

The root of the inheritance graph is a group of four predefined interfaces that make up the Reflective module; see Section 6.2, “The Reflective Interfaces,” on page 6-3. These interfaces collectively provide:

- operations that implement meta-object identity,
- operations for finding a meta-object’s containing package instance(s),
- an operation for finding a meta-object’s M2-level description, and
- operations for exercising the functionality of a meta-object independent of its generated interface.

Note – The interfaces in the Reflective module are all designed to be “abstract”; i.e. it is not anticipated that they should be the “most derived” type of any meta-object.

The interfaces for the Package objects, Association objects, Class Proxy objects and Instance objects provide functionality as described previously. The inheritance patterns are as follows:

- All Package object interfaces inherit (directly or indirectly) from RefPackage.
- All Association object interfaces inherit from RefAssociation.
- All Class Proxy interfaces inherit (directly or indirectly) from RefObject.
- All Instance interfaces inherit from the corresponding Class Proxy interfaces.
- When an M2-level Package P2 inherits from another P1, the corresponding interface P2 inherits from P1.
- When an M2-level Class C2 inherits from another C1:
 - the Class Proxy interface for C2 inherits from the Class Proxy for C1, and
 - the Instance interface for C2 inherits from the Instance for C1.

The diamond pattern of interface inheritance is virtually unavoidable. The C2's Class Proxy needs to inherit the interface features for C1's classifier-scoped Attributes and Operations. Similarly, C2's Instance interface needs to inherit the instance-scoped interface features.

Note – The IDL mapping supports some Tags for specifying additional IDL supertypes of various generated interfaces; see Section 5.6.3, “Tags for specifying IDL inheritance,” on page 5-36. The effect of these Tags on the inheritance graph is defined by the relevant IDL templates; see sections Chapter 5.8.4, Chapter 5.8.6, Chapter 5.8.7 and Chapter 5.8.10.

5.3 Computational Semantics for the IDL Mapping

This section specializes the MOF's general computational semantics (see Chapter 4, “The MOF Abstract Mapping”) for the MOF to IDL mapping.

5.3.1 Value Types and Equality in the IDL Mapping

The IDL mapping defines all MOF Instance types as CORBA object types that are descended from the “RefObject” interface; see Section 6.2.3, “Reflective::RefObject,” on page 6-10. Equality of Instance objects should be implemented as follows:

- Existing Instance objects are equal if and only if the “refMofId” operation defined by Section 6.2.2, “Reflective::RefBaseObject,” on page 6-5 returns the same string for both objects.
- Non-existent Instance objects are deemed to be equal if and only if they have the same object reference; i.e. when the “Object::_is_equivalent” operation returns true.

Note – An implementation must take care when comparing Instance object values to distinguish between non-existent (i.e., deleted) Instance objects and objects that may only be temporarily inaccessible. An operation should only raise an exception for a non-existent Instance object when it cannot be performed. In particular, an operation that replaces or removes defunct links or Instance values should not complain that the Instance being removed is defunct.

The MOF data types supported in the IDL mapping are the following CORBA data types:

- Primitive types - **boolean, char, octet, wchar, short, unsigned short, long, unsigned long, long long, unsigned long long, float, double, long double, fixed, strings, and wide strings.**
- Constructed types - arrays, sequences, enumerations and records.
- Type aliases - typedefs.

The native types supported in the IDL mapping are:

- CORBA union types.
- CORBA's **Any** and **TypeCode** data types.
- Ordinary CORBA object types (i.e., object types that are not descended from the “RefObject” interface).

Note – The IDL mapping currently does not support CORBA 2.3 value types, box types and abstract interface types.

Equality semantics for the standard MOF data values are as previously defined in Section 4.2.1, “Semantics of Equality for MOF Values,” on page 4-3. The standard rules are extended with the following rules for instances of IDL specific native types and CORBA object references:

- CORBA **TypeCode** values are equal if and only if they are equal according to the definition of `TypeCode::equal` in the CORBA Core specification.
- CORBA **Any** values are equal if and only they have equal types (according to the definition of `TypeCode::equal`), and their embedded values are equal according to the MOF definitions of data type equality.
- Ordinary CORBA object references and embedded object references for Instance objects are equal if and only if “`Object::_is_equivalent`” operation returns true.

The rule for CORBA object references applies both to “ordinary” object references and to object references for Instance objects. Similarly, it applies equally whether the object reference is the complete value, a component of a constructed value, or embedded within an **Any** value.

5.3.2 Lifecycle Semantics for the IDL Mapping

This section defines the IDL mapping's computational model for meta-object creation and deletion. It also gives definitions of copy semantics, though these should currently be viewed as indicative rather than normative.

5.3.2.1 Package object creation and deletion semantics

An M1-level Package object for a non-nested M2-level Package is created by invoking the create operation provided by the corresponding PackageFactory object. This create operation requires the caller to supply the values for all non-derived classifier-scoped Attributes. If the supplied initial values do not have the correct multiplicity or if they individually or collectively violate immediate Constraints defined in the metamodel, the create operation should raise an exception.

Instances of the following dependent M1-level objects are automatically created along with each M1-level Package object:

- An M1-level Package object is created for each nested Package within the outermost Package extent.
- An M1-level Package object is created for each clustered Package within the outermost Package extent.
- An M1-level Class Proxy object is created for each Class within the outermost Package extent.
- An M1-level Association object is created for each Association within the outermost Package extent.

The object references for the dependent Package and Class objects provide the “ref” attributes in the respective Package objects. The objects are initialized so that the `outermost_package` and `enclosing_package` operations return the appropriate M1-level Package objects.

Note – If an M2-level Package P2 clusters an existing top-level M2 Package P1, the above rules mean that two kinds of M1-level P1 Package objects can exist. If the user calls `create` on a P2 Package Factory object, the resulting P2 Package object will have its own dependent P1 Package object. On the other hand, if the user calls `create` on a P1 Package Factory, the resulting P1 Package object will be an outermost Package object. These two kinds of P1 Package objects behave identically, apart from their respective “`refOutermostPackage`” and “`refOutermostPackage`” operations; see Section 6.2.2, “`Reflective::RefBaseObject`,” on page 6-5.

When an M1-level Class Proxy objects is created, the values of the non-derived classifier-level Attributes are initialized from the corresponding create operation arguments. The “`all_of_type`” and “`all_of_kind`” collections will initially be empty, since no M1-level Instance objects will have been created in the Class Proxy extent.

Note – An implementation may support other mechanisms for creating or recreating outermost M1-level Package objects. Any such mechanism must also (re-)create and initialize the necessary dependent objects as above.

An outermost M1-level Package object can be destroyed using the “refDelete” operation; see Section 6.2.2, “Reflective::RefBaseObject,” on page 6-5. The required computational semantics for deleting an outermost Package object are straightforward. The following things must occur (in an unspecified order):

- The binding between the outermost Package object and its object reference(s) must be revoked.
- The bindings between all dependent Package, Association and Class Proxy objects and their object references must be revoked.
- All Instance objects within the extent of the outermost Package object must be destroyed as described below.

Note – A typical implementation will delete the metadata and reclaim the space used to store it. However, this behavior is not essential and in some situations it could be undesirable.

Dependent M1-level Package objects, M1-level Association objects and M1-level Class Proxy objects cannot be directly destroyed by the user. An implementation of the “refDelete” operation for these objects is required to raise an exception when called by client code. (The operations may be used to implement outermost Package deletion, but this is beyond the scope of this specification.)

5.3.2.2 *Instance object lifecycle semantics*

An M1-level Instance object can be created by invoking the appropriate create operation. Suitable create operations are present on both M1-level Class Proxy objects and M1-level Instance objects, depending on the M2-level Class inheritance graph. A create operation requires the caller to supply values for all non-derived instance-scoped Attributes for the Instance object. If any value does not conform to the Attribute’s multiplicity or if they individually or collectively violate any immediate Constraints on the meta-model, an exception is raised.

An Instance object is created within the extent of a Class Proxy object for the Instance’s M2-level Class. The Class Proxy can be found as follows:

1. Find the outermost Package extent containing the object on which the create operation was invoked.
2. Within that extent, find the one and only Class Proxy object for the M2 Class whose instance is being created.

If no Class Proxy can be found by the above, the create request violates the Supertype Closure Rule (see “The Supertype Closure Rule” on page 5-22) and an exception is raised.

Creation of an Instance object will also fail if the corresponding M2-level Class is abstract. Similarly, it will fail if the M2-level Class is a “singleton” Class and an Instance object for that Class already exists within the Class Proxy’s extent. In either case, an exception is raised.

When an Instance object is (successfully) created within the extent of a Class Proxy object, it becomes part of collection returned by the Class Proxy object’s “all_of_kind” operation. The Instance object remains a member of that collection for its lifetime; i.e. until it is deleted.

An Instance object will be deleted in the following three situations:

- when a client invokes the “refDelete” operation on the Instance object; see Section 6.2.3, “Reflective::RefObject,” on page 6-10,
- when the Package object for the Instance object’s outermost Package extent is deleted (see above), and
- when the Instance is a component of a “composite” Instance that is deleted. This applies to composites formed by both Associations and Attributes.

When an Instance object is deleted the following things must occur:

- The binding between the Instance object and its object reference(s) must be revoked.
- The Instance object must be removed from its Class Proxy object’s “all_of_type” collection.
- Any Instance objects that are components of the object being deleted must also be deleted.
- Links involving the deleted Instance object should be deleted as per the “Link lifecycle semantics” specification below.

An implementation will typically delete the state of an Instance object that has been deleted, and reclaim any associated space.

Note – When an Instance object is deleted, corresponding object reference values in non-composite Attributes of other objects become “dangling” references. These dangling references *should not* be automatically expunged or converted to nil object references, since doing so potentially destroys information and creates new structural errors. Instead, it is the user’s responsibility to ensure that dangling references in Attributes are tidied up in the most appropriate way.

5.3.2.3 *Link lifecycle semantics*

Links can be created and deleted in various ways. These include:

- by the user operations on M1-level Association objects; see Section 5.3.3, “Association Access and Update Semantics for the IDL Mapping,” on page 5-11,

- by the user operations corresponding to References on M1-level Instance objects; see Section 5.3.4, “Attribute Access and Update Semantics for the IDL Mapping,” on page 5-15.,
- by the user copying metadata (using some vendor specific API); see Section 4.10, “Recommended Copy Semantics,” on page 4-21,
- by the user deleting one or other linked Instance objects; see “Instance object lifecycle semantics” on page 5-9, and
- when the server notices that a linked Instance object no longer exists.

A link is created within the extent of an Association object, and becomes part of the collection returned by the Association object’s “links()” operation. A link remains within the extent in which it was created for the lifetime of the link; i.e. until it is deleted. When a link is deleted, it is removed from the “links” collection. Removing a link does not affect the lifecycle of the linked Instance objects.

According to Section 4.7.2.2, “Characteristics of M1-level Associations.,” on page 4-16, deletion of an Instance object causes any links for that object to become meaningless. Ideally, a well-formed M1-level Association instance should not contain such links. In practice, the immediate removal of meaningless links from an M1-level Association instance cannot always be implemented, in particular in the case of links that cross outermost Package extent boundaries.

Instead, a meta-object server is required to behave as follows. When an Instance object is deleted:

- all links referring to the Instance object that belong to Association instances within the same outermost Package extent as the Instance object *must* also be deleted, and
- any links referring to the Instance object that belong to Association instances in another outermost Package extent as the Instance object *may* also be deleted.

Note – The above semantics means that an Association instance can legally contain links that refer to defunct Instance objects in other extents.

5.3.3 Association Access and Update Semantics for the IDL Mapping

This section describes the computational semantics of the Association object access and update operations defined in the MOF to IDL Mapping and the Reflective interfaces. With a couple of exceptions, these semantics transform one *Well-formed State* (as defined in Section 4.7.2.1, “A Mathematical Model of Association State,” on page 4-15) to another. The exceptions are as follows:

- Deletion of an Instance object in another outermost Package extent may cause an Association instance to contain links that are not members of *Valid_Links*.
- Deletion of an Instance object can cause an *End_Links* set to contain fewer links than is required.

M1-level Instance objects are passed as CORBA object reference values in IDL mapped operations. However, since the Association State model requires that Links connect Instances, it is not legal to pass the CORBA nil object reference value as a parameter to any operation on an M1-level Association.

Note – While the semantics of Associations are described (below) in terms of sets of pairs of M1-level Instance objects, this should not be read as implying any particular implementation approach.

5.3.3.1 Access Operations

There are three kinds of link access operations in the M1-level Association interface generated by the IDL mapping:

- The “all_links” operation returns the current *Link_Set* for an Association object.
- The “<end_name>” operations return a projection of the corresponding *End_Links* sets.
- The “exists” operation tests for the existence of a given *Link* in the *Link_Set*.

These operations are defined to be side-effect free; i.e. they do not modify the *State* of the Association instance.

5.3.3.2 Link Addition Operations

The operations for adding links to an M1-level Association vary, depend on whether it has an ordered M2-level AssociationEnd:

- For an unordered Association, the “add” operation adds a *Link* to the *Link_Set*.
- For an ordered Association, the “add” and “add_before” operations both add a *Link* between a pair of Instances to the *Link_Set*. In the “add” case, the new *Link* is added after existing Links. In the “add_before” case, the new *Link* is added immediately before the link selected by the “before” argument.

More precisely, assuming that the first AssociationEnd is the ordered one and the new *Link* connects Instances *i* and *j*. The *Before* mapping is updated as follows:

- For “add”, all *Links* that were in *End2_Links_j* prior to the operation are *Before* the new *Link* when it completes.
- For “add_before”, the *Before_Link* connects the “before” and *j* Instances. For all *Links* that were in *End2_Links_j* and were *Before* the *Before_Link* prior to the operation, the pre-existing *Link* is *Before* the new *Link* after the operation. For all other *Links* that were in *End2_Links_j* prior to the operation, the new *Link* is *Before* the pre-existing *Link* after the operation.
- In both cases, the ordering of the other *End2_Links* sets are unchanged.

A number of constraints apply to the link addition operations:

- A new *Link* can only be added between extant Instances; i.e. the new *Link* must be a member of *Valid_Links*.

- An operation cannot add a *Link* that is already a member of the *Link_Set*.
- An operation cannot add a *Link* if it would make the number of members of either *End1_Links_i* or *End2_Links_j* greater than the respective AssociationEnd's "upper" bound.
- An operation cannot add a *Link* that creates a Composition cycle, or that violates the Composition or Reference Closure rules.

5.3.3.3 *Link Modification Operations*

There are two "modify" operations for replacing an existing *Link* in the *Link_Set* of an M1-level Association. One operation (in effect) modifies the Instance at the first end of a *Link*, and the second modifies the Instance at the second end. While the operation signatures do not vary, the semantics of the "modify" operations depend on whether the M2-level Association has an ordered AssociationEnd.

- In the non-ordered case, a "modify" operation is almost identical to a "remove" operation followed by an "add" operation. The only difference is in the bounds checking; see below.
- In the ordered case, a "modify" operation can differ from an "add" followed by a "remove" in the way that the *Before* ordering is handled. Specifically, if we assume that the first AssociationEnd is the ordered one, the *Before* mapping is updated as follows:
 - For "modify_<end1_name>(i, j, k)", the new *Link* (between *k* and *j*) occupies the same position in the *Before* ordering of *End2_Links_j* as the *Link* (between *i* and *j*) that it replaces.
 - For "modify_<end2_name>(i, j, k)", the new *Link* (between *i* and *k*) becomes the last *Link* in the *Before* ordering of *End2_Links_k*.
 - In both cases, the ordering of the other *End2_Links* sets are unchanged.

A number of constraints apply to the link modification operations:

- The *Link* that is replaced by the "modify" operation must be a member of *Link_Set*. However, it need not be a member of *Valid_Links*.
- The replacement *Link* that is created by a "modify" operation must be a member of *Valid_Links*.
- The replacement *Link* cannot already be a member of the *Link_Set*.
- A "modify" operation cannot produce a *Link* that would make the number of members in either the *End1_Links_k* or *End2_Links_k* sets greater than the respective AssociationEnd's "upper" bound.
- A "modify" operation cannot remove a *Link* if doing so would make the number of members of *End1_Links_i* or *End2_Links_j* less than the respective AssociationEnd's "lower" bound. (However, a *Link* **can** be produced in this situation.)
- A "modify" operations cannot produce a *Link* that creates a Composition cycle, or that violates the Composition or Reference Closure rules.

Note – A modify operation of the form “modify_<end1_name>(i, j, i)” is treated as a “no-op.” In particular, it does not trigger checking of “lower” or “upper” bounds.

5.3.3.4 *Link Removal Operations*

The “remove” operation can be used to delete an exist *Link* (between *i* and *j*) from the *Link_Set* of an M1-level Association. The constraints that apply to the link removal operation are:

- The operation cannot remove a *Link* if doing so would make the number of members of *End1_Links_i* or *End2_Links_j* less than the respective AssociationEnd’s “lower” bound.
- The operation cannot remove a *Link* that is not a member of the *Link_Set*. However, it should succeed if the *Link* is a member of *Link_Set* but not of *Valid_Links*.

5.3.3.5 *Changeability, Navigability and Derivedness*

The operation descriptions given above assume that the AssociationEnds of the M2-level Association have been defined with “isChangeable” and “isNavigable” set to true. If this is not so, the main impact is that certain operations are suppressed:

- If an AssociationEnd of an Association is defined as non-changeable (i.e. when its “isChangeable” flag is set to false), the IDL mapping suppresses various link update operations. The “add,” “add_before,” and “remove” operations are suppressed if either AssociationEnd is non-changeable. Furthermore, the “modify_<end_name>” operation is suppressed for any AssociationEnd that is non-changeable, along with any related Reference-based operations.
- If an AssociationEnd of an Association is defined as non-navigable (i.e. when its “isNavigable” flag is set to false) the IDL mapping suppresses any link operations that depend on the ability to search based on that AssociationEnd. Specifically, it suppresses the “<assoc_end>”, “add_before_<end>”, “modify_<end>” operations.

Setting “isDerived” to be true for an M2-level Association is a “hint” that an M1-level Association’s *Link_Set* and *Before* mapping should be computed from other M1-level information. Apart from this, the IDL mapping makes no distinction between derived and non-derived Associations. Equivalent IDL interfaces are generated in each case, and the semantics are defined to be equivalent. If a derived Association’s operations are coded by hand, it is the programmer’s responsibility to ensure that they implement the required semantics.

Some combinations of the Association and AssociationEnd flags result in generated interfaces that are of little use. For example:

- Setting “isChangeable” to be false on one AssociationEnd and not the other results in an M1-level Association that supports one “modify” operation but no “add” or “remove” operations.
- Setting “isChangeable” to be false on an Association which has “isDerived” set to false results in a “stored” Association with no operations to update the *Link_Set*.

5.3.4 Attribute Access and Update Semantics for the IDL Mapping

The IDL mapping maps M2-level Attributes to a variety of operations, depending on the Attribute's "multiplicity" settings. There are three major cases:

- single-valued with bounds of [1..1]),
- optional with bounds of [0..1], and
- multi-valued.

Unlike Associations, the CORBA "nil" object reference is a legal (and logically distinct) value for any Class or object reference-valued Attribute. When an accessor operation returns a "nil" object reference, this does not necessarily mean that the Attribute has no value(s). In addition, the lifecycle semantics for Attributes in the IDL mapping mean that an accessor operation can return a reference for a non-existent object.

Note – While the semantics of Attributes are described (below) in terms of notional relations between M1-level values, this should not be read as implying any particular implementation approach.

5.3.4.1 Single-valued Attributes

The interfaces and semantics for single-valued Attributes are the simplest to describe. A single-valued Attribute (i.e. one whose "lower" and "upper" bounds are set to one) is mapped to two IDL operations: 1) "<attr_name>" and 2) "set_<attr_name>."

The "<attr_name>" operation returns the current value of the named Attribute for an M1-level Instance object. In the single-valued case, this is a single Instance of the Attribute's M1-level base type as mapped by the IDL mapping. In the terminology of Section 4.4.1, "Attribute name and type," on page 4-4, the operation returns the M1-level value that is related to the Instance object by the notional "<attr_name>" Class — AttrType relation.

The "set_<attr_name>" operation replaces the current value of the named Attribute for an M1-level Instance with a new value. As before, the new value is a single Instance of the Attribute's M1-level base type as mapped by the IDL mapping. The operation replaces the existing Class — AttrType relationship with a new one between the Instance object and the new value.

The behaviour of "set_<attr_name>" for an Class-valued Attribute (i.e., one with "composite" aggregation semantics) is constrained as follows:

- The new value supplied must be either a reference to an existing Instance object or a nil object reference.
- The new value (i.e., the component Instance) must not already be a component of another Instance object.
- The composite and component Instance objects must belong to the same outermost M1-level Package extent (i.e., the Composition Closure rule must not be violated).

- Creating the new Class — AttrType relationship must not create a composition cycle.

5.3.4.2 *Optional Attributes*

The interfaces and semantics for optional Attributes are also relatively straightforward. An optional Attribute (i.e., one whose “lower” bound is 0 and whose “upper” bound is 1) maps to three operations:

- “<attr_name>”
- “set_<attr_name>”
- “unset_<attr_name>”

The IDL mapping treats an M1-level optional Attribute as having two states. In the “set” state, the Attribute has a value which is an instance of the Attribute’s M1-level base type. In the “unset” state, the Attribute has no value.

In the single-valued case, “<attr_name>” simply returns the current M1-level value for the Attribute. In the optional case, the semantics depend on whether the Attribute is currently “set” or “unset.”

- If the Attribute is “set” (i.e., there is a Class — AttrType relationship between the Instance object and some other value), the “<attr_name>” operation returns the related value.
- If the Attribute is “unset” (i.e., there is no Class — AttrType relationship with the Instance object in the “class” role), the “<attr_name>” operation raises an exception.

The “set_<attr_name>” operation behaves exactly as in the single-valued case; it replaces the existing Class — AttrType relationship (if any) with a relationship with the new value. As a consequence, the Attribute enters the “set” state. The structural constraints for “set_<attr_name>” in the single-valued case apply here as well.

The “unset_<attr_name>” operation removes the Class — AttrType relationship, if it exists, leaving the Attribute in the “unset” state.

5.3.4.3 *Multi-valued Attributes*

The interfaces and semantics for multi-valued Attributes are relatively complicated, and depend to a considerable extent on the settings of the “isOrdered” and “isUnique” fields of the M2-level Attribute’s “multiplicity” property.

M1-level operations on multi-valued Attributes can be divided into two groups. The “<attr_name>” and “set_<attr_name>” operations access and update the Attribute’s state as a single value, transferring it as a CORBA sequence type. The other operations treat the Attribute’s state as a collection of values, and update it by adding, modifying or removing individual elements of the collection.

The “<attr_name>” and “set_<attr_name>” operations transfer an Attribute’s M1-level state using a “collection” type. This is a named IDL sequence type whose base type is the Attribute’s M1-level base type, and whose name is determined by the “name” of the Attribute’s “type” and the settings of the “isOrdered” and “isUnique” flags; for details, see Section 5.7.1.6, “Literal String Values,” on page 5-41.

The “<attr_name>” operation returns the multi-valued Attribute’s value as a sequence using the IDL type described above. The contents of the result comprise the collection of base type instances related to the Instance object by the Class — AttrType relation. If “isOrdered” is true, the order of the Class — AttrType relationships determines the order of the elements in the sequence. If the collection is empty, the returned value is a zero length sequence.

The “set_<attr_name>” operation replaces the multi-valued Attribute’s value with a new collection of base type instances. If the Attribute is ordered, the order of the elements in the parameter value determines the order of the new Class — AttrType relationships.

A number of restrictions apply to the “set_<attr_name>” operation for multi-valued Attributes. These are as follows:

- If the Attribute’s “multiplicity” has the “isUnique” flag set to true, no two base type instances in the collection may be equal.
- If the Attribute’s “multiplicity” has a “lower” value greater than zero, there must be at least that many elements in the collection.
- If the Attribute’s “multiplicity” has a “upper” value other than the “UNBOUNDED” value (i.e., -1), there can be at most that many elements in the collection.

If the Attribute has composite semantics (i.e., the Attribute’s “type” is expressed using a Class) the following restrictions also apply:

- Each element (i.e., Instance object) in the new value collection must be either a reference to an existing Instance object or a nil object reference.
- No element of the new value collection can already be a component of another Instance object.
- The composite and every component Instance objects must belong to the same outermost M1-level Package extent (i.e., the Composition Closure rule must not be violated).
- Creating the new Class — AttrType relationships must not create any composition cycles.

The IDL mapping can define up to 7 additional operations for a multi-valued Attribute. There are up to 3 operations for adding new element values to an Attribute collection, up to 2 for modifying them and up to 2 for removing them. The subset that is available for a given Attribute depends on the “isUnique” and “isOrdered” flags in the M2-level Attribute’s “multiplicity.” This is shown in Figure 5-1 on page 5-2.

Table 5-1 Element Update Operations for Multi-valued Attributes

isOrdered	isUnique	Operations available
false	false	add_<attr_name>, modify_<attr_name>, remove_<attr_name>
false	true	add_<attr_name>, modify_<attr_name>, remove_<attr_name>
true	false	add_<attr_name>, add_<attr_name>_before, add_<attr_name>_at, modify_<attr_name>, modify_<attr_name>_at, remove_<attr_name>, remove_<attr_name>_at
true	true	add_<attr_name>, add_<attr_name>_before, modify_<attr_name>, remove_<attr_name>

When “isOrdered” is set to false, the operations provided are the basic ones for adding, modifying, or removing element values. Given that the collection is unordered, there is no need to specify the position at which a new element value is added, or (in the false, false case) which of a number of equal element values should be modified or removed. The semantics of the operations for an unordered Attribute are as follows:

- The “add_<attr_name>” operation creates a new Class — AttrType relationship between the Instance object and the M1-level base type instance being added to the Attribute collection.
- The “modify_<attr_name>” operation replaces the Class — AttrType relationship between the Instance object and the M1-level base type instance being modified with another for the new element value.
- The “remove_<attr_name>” operation removes the Class — AttrType relationship between the Instance object and the M1-level base type instance being removed from the Attribute collection. Removing the instance decreases the Attribute collection’s length rather than leaving a “hole.”

These three operations must also respect the restrictions listed above for the multi-valued “set_<attr_name>” operation.

When “isOrdered” is set to true, the “add_<attr_name>,” “modify_<attr_name>,” and “remove_<attr_name>” operations take on additional semantics:

- The “add_<attr_name>” operation must ensure that the newly added element appears as the last element in the Attribute collection.
- The “modify_<attr_name>” operation must ensure that the replacement M1-level base type instance appears in the same position in the Attribute collection as the value that it replaces. When “isUnique” is set to false, the collection may contain duplicates. In this case, the operation should replace the first example of the instance in the ordered Attribute collection.
- When “isUnique” is set to false, the “remove_<attr_name>” operation should remove the first example of the instance in the ordered Attribute collection.

In addition, the client is provided with extra operations for order sensitive element update:

- The “add_<attr_name>_before” operation is similar to the “add_<attr_name>” operation, except that the new instance is added to the Attribute collection before an existing element designated by the caller. When “isUnique” is false, the operation is defined to replace the first example of the instance in the Attribute collection.
- When “isOrdered” is true and “isUnique” is false, the “add_<attr_name>_at,” “modify_<attr_name>_at,” and “remove_<attr_name>_at” are provided to allow the client to update the collection in the presence of duplicates. These operations specify an element insertion point or an element to be modified to be removed by giving an position index. For the purposes of these operations, the elements in an Attribute collection are numbered starting from zero according to the defined order of the members of the collection. The operations are as follows:
 - **add_<attr_name>_at** - inserts the new M1-level base type instance so that it appears at the position given. The instance originally at that position, and all instances will have their position indexes increased by one.
 - **modify_<attr_name>_at** - replaces the M1-level base type instance at the position.
 - **remove_<attr_name>_at** - removes the M1-level base type instance at the position given. Any instances in the collection that follow the removed instance will have their position indexes decreased by one (i.e., the operation does not leave a “hole” in the Attribute collection).

These five additional operations must also respect the restrictions listed above for the multi-valued “set_<attr_name>” operation.

5.3.4.4 *Changeability and Derivedness*

The previous semantic descriptions assume the M2-level Attribute has “isChangeable” set to true and “isDerived” set to false. This subsection describes what happens if this is not the case.

If an Attribute has “isChangeable” set to false, the effect on the IDL mapping is that all generated operations for updating the Attribute’s state are suppressed. This does not preclude the existence of other mechanisms for updating the Attribute’s state.

Setting an Attribute’s “isDerived” flag to true, has no effect on the IDL mapping. The operations generated for the derived and non-derived cases are equivalent and they are defined to have equivalent semantics. If a derived Attribute’s operations are coded by hand, it is the programmer’s responsibility to ensure that they implement the required semantics.

5.3.4.5 *Classifier scoped Attributes*

The previous semantic descriptions assume the M2-level Attribute has “scope” set to “instance_level.” When an Attribute’s “scope” is “classifier_level,” we can model the notional relation that defines the M1-level Attribute state as a relation between the

Class extent and the AttrType; see Section 4.4.3, “Scope,” on page 4-6. In the IDL mapping, this translates to a notional relation between a Class Proxy object and instances of the Attribute’s M1-level base type.

On this basis, an Attribute whose “scope” is “classifier_level” differs from one whose “scope” is “instance_level” in the following respects:

- The notional Class Proxy — AttrType relation supplies the value or values accessed and updated by “classifier_level” scoped Attribute operations.
- When the Attribute has aggregation semantics of “composite”:
 - the Composition Closure rule means that the Class Proxy object and M1-level Attribute value Instances must belong to the same extent, and
 - checking for composition cycles is unnecessary. The Class Proxy object that holds the Attribute value(s) is not an Instance, and thus cannot be a “component” in this sense.

5.3.4.6 *Inherited Attributes*

The previous semantic descriptions apply equally to Attributes defined within an M2-level Class, and Attribute inherited from supertypes of the Class.

5.3.4.7 *Life-cycle Semantics for Attributes*

The previous semantic descriptions say nothing about how an Attribute gets its initial value or values. (With the exception of the single-valued case of the “<attr_name>” operation, the semantic descriptions would “work” if no notional relationships existed initially.) In fact, the IDL mapping ensures that all M1-level Attributes get a client-supplied initial value:

- All “instance_level” scoped Attribute values for a M1-level Instance object are initialised from the parameters to the “create_<class_name>” operation.
- All “classifier_level” scoped Attribute values within the extent of an outermost M1-level Package are initialised from the parameters to the “create_<package_name>” operation.

An M1-level Attribute only exists while the M1-level Instance object or Class Proxy object that it belongs to exists. When the object is deleted, the notional relationships disappear as well.

Attributes with “composite” aggregation semantics have special life-cycle. When an object with a composite Attribute is deleted, the Instance object or objects that form its value are also deleted.

Note that unlike Associations, when an Instance object is deleted, the delete operation should make no attempt to tidy up “dangling references” to it.

5.3.5 Reference Semantics for the IDL Mapping

The IDL mapping maps References into a hybrid that combines an Attribute style interface with Association access and update semantics. In each case, a Reference operation maps fairly directly onto an Association operation as shown in Figure 5-2 below.

Table 5-2 Semantic mapping of Reference operations to Association Operations

Multiplicity	Reference Operation	Association Operation(s) (assuming that the referenced AssociationEnd is the 2nd one)
optional	<code>i.<reference_name>()</code>	<code>temp = a.<referenced_end_name>(i)</code> <code>if temp.size > 0 then</code> <code>temp[0]</code> <code>else</code> <code>raise NotSet</code>
single- and multi-valued	<code>i.<reference_name>()</code>	<code>a.<referenced_end_name>(i)</code>
optional	<code>i.set_<reference_name>(new)</code>	<code>old = a.<reference_end_name>(i)</code> <code>if old.size > 0 then</code> <code>a.modify_<reference_end_name>(i, old[0], new)</code> <code>else</code> <code>a.add(i, new)</code>
optional	<code>i.unset_<reference_name>()</code>	<code>old = a.<reference_end_name>(i)</code> <code>if old.size > 0 then</code> <code>a.remove(i, old[0])</code>
single-valued	<code>i.set_<reference_name>(new)</code>	<code>old = a.<ref_end_name>(i)</code> <code>a.modify_<ref_end_name>(i, old, new)</code>
multi-valued	<code>i.set_<reference_name>(new)</code>	<code>old = a.<ref_end_name>(i)</code> <code>for j in 0 .. (old.size - 1) do</code> <code>a.remove(i, old[j])</code> <code>for j in 0 .. (old.size - 1) do</code> <code>a.add(i, new[j])</code>
multi-valued	<code>i.add_<reference_name>(new)</code>	<code>a.add(i, new)</code>
multi-valued	<code>i.add_before_<reference_name>(new, before)</code>	<code>a.add_before_<referenced_end_name>(i, new, before)</code>
multi-valued	<code>i.modify_<reference_name>(old, new)</code>	<code>a.modify_<referenced_end_name>(i, old, new)</code>
multi-valued	<code>i.remove_<reference_name>(old)</code>	<code>a.remove_<referenced_end_name>(i, old)</code>

In practice, an implementation also needs to transform exceptions reported for the Association operations into exceptions that apply from the Reference perspective. In addition, a “quality” implementation would ensure that Reference operations did not leave the Association object in a half way state following an exception.

Note – The above semantic mapping description is not intended as implying any particular implementation approach.

5.3.6 Cluster Semantics for the IDL Mapping

The impact of clusters on the IDL mapping semantics are largely described elsewhere. At the M1-level, a clustered Package behaves identically to a nested Package in terms of life-cycle and extent rules. The only significant difference is that clustering is not always a strict composition relationship at the M1-level; see Section 4.6.4, “Package Extents,” on page 4-11. In the IDL mapping, this means that two or more Package “ref” attributes to point at the same clustered Package instance.

5.3.7 Atomicity Semantics for the IDL Mapping

All operations defined by the IDL mapping (including the Reflective versions) are required to be atomic and idempotent:

- If an operation succeeds, state changes required by the specification should be made, except as noted below:
 - When an Instance object is deleted, deletion of any component Instance objects may occur asynchronously.
 - When an Instance object is deleted, removal of links to the deleted Instance object may occur asynchronously.
- If an operation fails (e.g., by raising an exception), no externally visible changes should be caused by the failed operation.
- When the invocation of two or more operations overlaps in time, the resultant behaviour should be semantically equivalent to the sequential invocation of the operations in some order.

Note – The IDL mapping specification does not require a transactional or persistent implementation of a meta-data server.

5.3.8 The Supertype Closure Rule

The inheritance pattern for Instance and Class Proxy interfaces has an important consequence when one M2-level Class is a sub-Class of a second one.

Recall that each Class Proxy interface defines a factory operation for the corresponding Instance object, and that it also inherits from the Class Proxy interfaces for any M2-level super-Classes. Taken together, this means that any Class Proxy object has operations for creating Instance objects for both the M2-level Class, and all of its M2-level super-Classes.

Normally, this artifact of the IDL inheritance hierarchy is just a convenience. However, problems arise when an M2-level Class (e.g., P2::C2) has a super-Class that is imported from another M2-level Package (e.g., P1::C1); see Figure 5-3 on page 5-23.

The Class Proxy interface corresponding to the C2 Class now has a factory operation to create instances of a Class from another Package, and therefore would appear to require all of the mechanisms for creating, accessing, updating and deleting these instances. This is not what Package importing is defined to mean.

The adopted solution to this problem is to add an extra restriction to the MOF computational semantics. This restriction is known as the *Supertype Closure Rule*.

Supertype Closure Rule

Suppose that the Package extent for a non-nested M2-level Package P contains a Class Proxy object which has a create operation for instances of Class C. This create operation can be used if and only if the M2-level closure of the Package P under generalization and clustering includes the M2-level Class C.

In other words, a factory operation for instances of an M2-level Class will only work within a Package instance with the machinery for supporting the Class. The Supertype Closure Rule is illustrated in Figure 5-3.

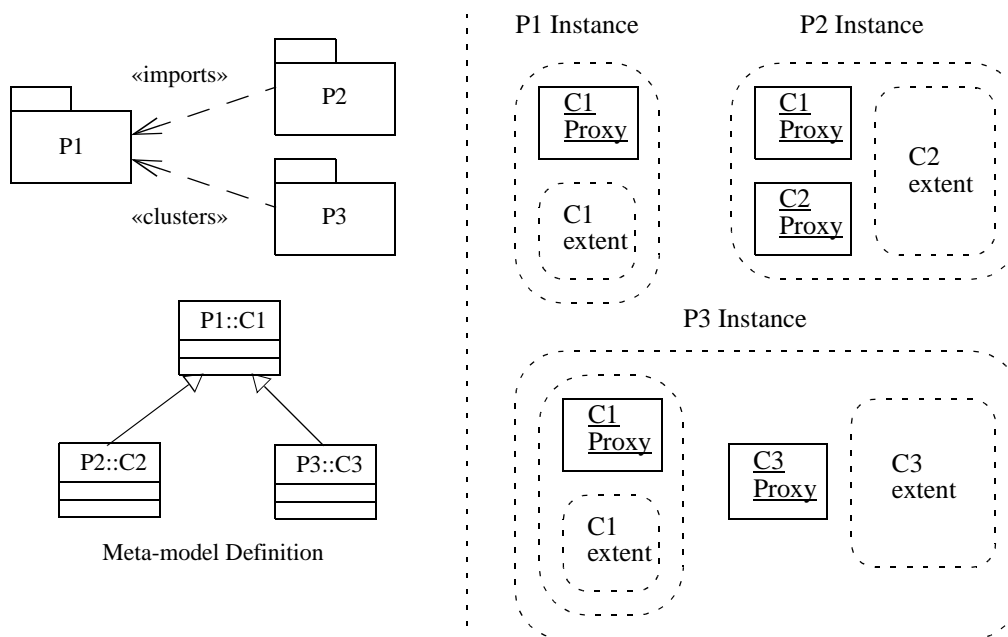


Figure 5-3 Supertype Closure Rule

5.3.9 Copy Semantics for the IDL Mapping

The IDL mapping currently defines no APIs for copying meta-data. Copy semantics are therefore beyond the scope of this chapter.

5.4 Exception Framework

This section describes the way that Exceptions are organized in the MOF to IDL mapping. These exceptions are raised in a variety of CORBA interfaces, including:

- Reflective interfaces: (see Section 6.2.2, “Reflective::RefBaseObject,” on page 6-5, Section 6.2.3, “Reflective::RefObject,” on page 6-10, Section 6.2.4, “Reflective::RefAssociation,” on page 6-25 and Section 6.2.5, “Reflective::RefPackage,” on page 6-32).
- Model interfaces (see Section 3.4, “MOF Model Classes,” on page 3-15 and Section 3.5, “MOF Model Associations,” on page 3-72).
- Specific interfaces produced by the mapping templates (see Section 5.8, “IDL Mapping Templates,” on page 5-46).

The exceptional conditions that arise in the context of the MOF to IDL mapping are classified into 5 groups:

1. Structural errors - this group covers those situations where the basic structural consistency rules for the metadata are (or would be) violated. For example, when there are too many or too few elements in a collection value.
2. Constraint errors - this group covers violations of metadata consistency rules specified in the metamodel using Constraints.
3. Usage errors - this group covers those situations where a client tries to use the MOF interfaces in a meaningless way. For example, giving a ‘position’ for a collection element that is outside of the collection bounds.
4. Reflective errors - this group covers errors that can only occur when using the Reflective interfaces. For example, calling “refInvokeOperation” on an Attribute. These errors are the notional equivalent of runtime type errors.
5. Semantic errors - this group covers errors not covered above (i.e., implementation specific errors).

The complexity of the MOF means that the number of exceptional conditions is (at least in theory) unbounded. The precise set of possible exceptional conditions for just one operation in the mapped interfaces can be very hard to define. Constraint and Semantic errors are particularly difficult to tie down. Furthermore, including lots of exceptions in an IDL operation signature can make client code inordinately complex.

To solve these problems, the MOF IDL mapping defines the `MofError` exception that covers most of the exceptional conditions that might arise.

```

struct NamedValueType {
    string name;
    ValueType value;
};
typedef sequence < NamedValueType > NamedValueList;
exception MofError {
    string error_kind;
    RefBaseObject element_in_error;

```

```

NamedValueList extra_info;
string error_description;
};

```

The fields of the `MofError` exception are defined as follows:

- **error_kind** is a string that denotes the particular kind of exceptional condition that is being raised. The formation of values for this field is discussed below.
- **element_in_error** is the `DesignatorType` for the object or feature that is deemed to be in error for this error condition. The detailed specifications of the error conditions below define which meta-object should be returned in each case. In situations where no M2-level meta-objects are available, this field may contain a nil object reference.
- **extra_info** is a list of name / value pairs that provides the client with extra information about the error condition.

The list consists of zero or more standardized name / value pairs, followed by any implementation specific pairs. For the standardized part of the list, the sequence of the pairs and the values (including casing) of the names are mandatory. This allows clients to extract list elements by position or by matching names. It is recommended that implementers take the same approach for the implementation specific part of the list.

- **error_description** is a human readable diagnostic message. The contents of this field are not specified by this document.

Note – The standardized name / value pairs for the **extra_info** field represent a compromise between the anticipated cost of implementation and the provision of useful information to the caller. Implementers are encouraged to provide additional information. Similarly, implementers are encouraged to provide detailed and informative diagnostics in the **error_description** field.

5.4.1 *Error_kind string values*

The values of the **error_kind** field or `MofError` are structured using Java's reversed domain name syntax:

```

"org.omg.mof:structural.composition_cycle"
"au.edu.dstc.mofamatic:botched_assertion"

```

The values for each group of errors are as follows:

- Structural and Reflective errors: the prefix **"org.omg.mof:"** followed by either **"structural."** or **"reflective."** and then the specific error name in lowercase with underscores between words. These values are defined as constants in the IDL for the Reflective module.

- Constraint errors: the IDL prefix for the metamodel (if any), followed by “:**constraint.**” followed by the qualified constraint name using the Format2 convention. For example, a Constraint named “MyConstraint” declared in “PackageA::ClassB,” the error kind string value is:

“:**constraint.package_a.class_b.my_constraint**”

or with an IDL prefix of “**com.acme**”, it is:

“**com.acme:constraint.package_a.class_b.my_constraint**”

See Section 5.8.17, “Constraint Template,” on page 5-91 for the definitive specification.

- Usage errors: not applicable. None of these error conditions are signalled using MofError.
- Semantic errors: an implementation specific prefix, followed by “:**semantic.**” followed by an implementation specific string. It is strongly recommended that the implementation specific part follow the conventions above (i.e., reverse domain names, all lowercase, periods for qualification and underscores between words).

5.4.2 Structural Errors

All structural errors are signalled using MofError. With the exception of “Underflow,” the consistency rules covered by the structural errors are either pre- or post-conditions on operations.

The MOF IDL mapping defines the structural errors as defined in Table 5-3.

Table 5-3 Structural Errors signalled using MofError

Structural error	“Element_in_error”	Standard “extra_info”	Description
Underflow	Attribute, Parameter, or Association End defining the Multiplicity that is violated.	none	<p>“Underflow” arises when a collection or projection contains fewer values than is required by the corresponding Multiplicity.lower.</p> <p>Note that the evaluation “underflow” is context dependent. For an operation which takes a collection value as a parameter, or whose net effect is to decrease the number of elements in a multi-valued Attribute or a projection of a Link set, “underflow” is treated as an immediate constraint. In other cases, “underflow” is treated as a deferred constraint.</p>

Table 5-3 Structural Errors signalled using MofError (continued)

Structural error	“Element_in_error”	Standard “extra_info”	Description
Overflow	Attribute, Parameter, or Association End defining the Multiplicity that is violated.	none	"Overflow" arises when a collection or projection contains more values than is allowed by the corresponding Multiplicity.upper.
Duplicate	Attribute, Parameter, or Association End defining the Multiplicity that is violated.	“duplicate” : Any(<Value>) A value that appears more than once in the unique collection / projection.	"Duplicate" arises when a collection or projection whose corresponding Multiplicity.is_unique is true contains duplicate values. For example, when two or more values at different positions in the collection or projection that are “equal” according to the definitions in Section 4.2.1, “Semantics of Equality for MOF Values,” on page 4-3.
Reference Closure	Reference for which the closure rule is violated.	“external” : Any(<Instance>) An Instance that violates a closure rule with respect to the Association being updated.	"Reference Closure" can arise when an Association extent contains a link for an Instance object belonging to another outermost Package extent. More particularly, this happens when the Instance object's M2-level Class (or a super-Class ancestor) has a Reference to the M2-level Association. See Section 4.9.1, “The Reference Closure Rule,” on page 4-19.
Composition Closure	Attribute or Association for which the closure rule is violated.	“external” : Any(<Instance>) An Instance that was passed as or within in an operation parameter that violates the closure rule.	"Composition Closure" arises when an Instance object is member of a composite which crosses an outermost Package extent boundary. See Section 4.9.2, “The Composition Closure Rule,” on page 4-20
Supertype Closure	Class of the object that cannot be created.	none	"Supertype Closure" arises when a client attempts to create an Instance object in a Package extent that does not support its M2-level Class. See Section 5.3.8, “The Supertype Closure Rule,” on page 5-22.
Composition Cycle	Attribute, Reference or Association that is being updated to form the cycle..	“cyclic” : Any(<Instance>) A composite Instance passed as or within a parameter that would become cyclic as a result of this operation.	"Composition Cycle" arises when an Instance object is a component of itself via one or more relationships defined by composite Associations or composite Attributes.
Nil Object	Reference or Association End for which the nil object reference was supplied.	none	"Nil Object" arises when a Association operation is passed a CORBA nil object reference.

Table 5-3 Structural Errors signalled using MofError (continued)

Structural error	“Element_in_error”	Standard “extra_info”	Description
Inaccessible Object	Attribute, Parameter, Reference or Association End for which the inaccessible object was detected.	“inaccessible” : Any(<RefObject>) An Instance object that was inaccessible.	"Inaccessible Object" arises when an operation tries to use an Instance object only to find that it is currently inaccessible.
Invalid Object	Attribute, Parameter, Reference or Association End for which the invalid object was detected.	“invalid” : Any(<RefBaseObject>) An object reference for a MOF meta-object that does not exist.	"Invalid Object" can arise when an object operation detects a reference for a non-existent (i.e.,deleted) object.
Already Exists	Class of the object that already exists	“existing” : Any(<Instance>) The pre-existing singleton Instance object for the extent.	“Already Exists” arises when a client attempts create a second Instance object for an M2-level Class with “isSingleton” of true.

Note – There are no mandatory ‘extra_info’ pairs for "Overflow" and "Underflow" because the error conditions occur in such a wide range of contexts that it is difficult to come up with a set that is universally applicable. Vendors are encouraged to innovate by defining non-standard pairs.

The following IDL constants define the corresponding **error_kind** strings.

```

const string UNDERFLOW_VIOLATION =
    "org.omg.mof:structural.underflow";
const string OVERFLOW_VIOLATION =
    "org.omg.mof:structural.overflow";
const string DUPLICATE_VIOLATION =
    "org.omg.mof:structural.duplicate";
const string REFERENCE_CLOSURE_VIOLATION =
    "org.omg.mof:structural.reference_closure";
const string SUPERTYPE_CLOSURE_VIOLATION =
    "org.omg.mof:structural.supertype_closure";
const string COMPOSITION_CYCLE_VIOLATION =
    "org.omg.mof:structural.composition_cycle";
const string COMPOSITION_CLOSURE_VIOLATION =
    "org.omg.mof:structural.composition_closure";
const string NIL_OBJECT_VIOLATION =
    "org.omg.mof:structural.nil_object";
const string INACCESSIBLE_OBJECT_VIOLATION =
    "org.omg.mof:structural.inaccessible_object";
const string INVALID_OBJECT_VIOLATION =
    "org.omg.mof:structural.invalid_object";
const string ALREADY_EXISTS_VIOLATION =
    "org.omg.mof:structural.already_exists";

```


5.4.3 Constraint Errors

Constraint errors occur when a consistency rule is defined as a Constraint in the metamodel.

All Constraint errors are signalled by raising **MofError**. The fields of the **MofError** exception are defined as follows:

- The **error_kind** string is defined by the IDL mapping rules (see Section 5.8.17, “Constraint Template,” on page 5-91).
- The **element_in_error** is the designator for the Constraint that has been violated.
- The value of the **extra_info** field is implementation specific. Where possible, the implementation should provide the constrained object(s) or value(s) for which the constraint is violated.

Constraints can be defined with an “evaluationPolicy” of “immediate” or “deferred.” In the former case, violations of the rule are likely to be reported when a constrained object is created or updated. In the latter case, violations are likely to be reported when deferred Constraint checking is triggered.

Note – The above statements assume that constraint checking is implemented according to the spirit of Section 4.6, “Extents,” on page 4-9.

5.4.4 Semantic Errors

The Semantic error group is the “catch all” for otherwise unclassified implementation specific errors. Semantic errors are signaled by raising the **MofError** exception with appropriate. Possible sources of this error include:

- additional metadata consistency rules that are not specified in the metamodel,
- implementation specific access control violations,
- resource limitations in a metadata server, and
- internal errors in a metadata server.

The values of the **MofError** exception fields for a Semantic error are implementation specific:

- Implementors should define a unique strings for the **error_kind** field to distinguish the different kinds of Semantic error. These values should conform to the pattern described in Section 5.4.1, “Error_kind string values,” on page 5-25.”
- The values and meanings of the **element_in_error** and **extra_info** fields should be defined as appropriate.

5.4.5 Usage Errors

The Usage error group indicate inappropriate use of the MOF IDL interfaces. They can arise when a client is using either the Reflective interfaces, or the interfaces generated by the IDL mapping.

The Usage errors are signalled using their own exceptions,

Table 5-4 Usage Exceptions

Usage Exception	Arguments	Description
NotFound	none	NotFound is raised by modify and remove operations on multi-valued Attributes, References and Associations when the argument that should identify the member or link to be removed does not match any value that is currently there.
NotSet	none	NotSet is raised when a client attempts to read the element value of an optional collection (i.e. one with bounds of [0..1]) when the collection is empty.
BadPosition	none	BadPosition is raised by a positional add, modify or remove operation is supplied with a 'position' argument whose value is out of range. The collection's current size is returned in the exception's 'current_size' field. This will be 0 if the collection is empty, 1 if it contains a single member and so on.

Note – The members of a collection value containing *size* elements are numbered {0, 1,... *size* - 1} for the purposes of the positional update operations. The positional modify / remove operations are defined to modify or remove the member indexed by the **position** (i.e., position values in the range 0 to *size* - 1 inclusive are valid). The positional add operation is defined to insert a member before the member indicated by the **position**. In this case, position values in the range 0 to *size* inclusive are valid, with *size* meaning "insert at the end."

The IDL declarations for the **Usage** error exception are as follows:

```

exception NotFound {};
exception NotSet {};
exception BadPosition {
    unsigned long current_size;
};

```

5.4.6 Reflective Errors

Reflective error conditions occur exclusively in operations in the Reflective interfaces. They occur when a Reflective operation is invoked with parameters that contradict the target object's description in the metamodel. When the client uses interfaces generated by the IDL mapping, the static type checking based on the specific IDL signatures should prevent the equivalent errors from occurring.

In most cases, the MofError exception is used to signal reflective errors. Table 5-5 lists the Reflective errors that are signalled using MofError, along with the MofError field specifications and descriptions. All are pre-conditions for the respective operations.

Table 5-5 Reflective Errors signalled using MofError

Reflective error	“Element_in_error”	Standard “extra_info”	Description
Invalid Designator	ModelElement that is invalid	none	"Invalid Designator" arises when a "feature" parameter: <ul style="list-style-type: none"> • is not a Model::ModelElement, or • does not denote an accessible CORBA object.
Wrong Designator Kind	ModelElement that has the wrong kind	none	"Wrong Designator Kind" arises when the supplied designator has an inappropriate most-derived type. For example, when a Model::Attribute is supplied where a Model::Operation is required.
Unknown Designator	ModelElement that is not known.	none	"Unknown Designator" arises when the supplied designator does not belong in this context. For example, when a Model::Attribute is not a member of this Instance's Class or its superClasses.
Abstract Class	Class that is abstract.	none	“Abstract Class” arises when a client calls “refCreateInstance” for a Class that is defined as abstract.
Not Changeable	ModelElement that has “isChangeable” = false	none	"Not Changeable" arises when an update operation is attempted on something that is defined by the metamodel to be not changeable.
Not Navigable	AssociationEnd that has “isNavigable” = false	none	"Not Navigable" arises when RefAssociation operations are attempted for an AssociationEnd that is defined by the meta-model to be not navigable.
Not Public	ModelElement that has “visibility” = "private_vis" or "protected_vis"	none	"Not Public" arises when an operation is attempted for a "private" or "protected" feature.
Wrong Scope	Attribute or Operation with “scope” = “instance_level”	none	"Wrong Scope" arises when an attempt is made to use an instance-level Attribute or Operation from a Class proxy object.

Table 5-5 Reflective Errors signalled using MofError

Reflective error	“Element_in_error”	Standard “extra_info”	Description
Wrong Multiplicity	Reference or Attribute used in error	none	<p>“Wrong Multiplicity” arises when a reflective operation is requested where the corresponding specific operation does not exist for this feature’s multiplicity. For example:</p> <ul style="list-style-type: none"> • a member update on a [0..1] or [1..1] feature, • a unset on a feature that is not [0..1], • an add_value_at on an unordered feature.
Wrong Type	Attribute, Reference, AssociationEnd, or Parameter for the value that is in error.	<p>“invalid_value” : Any</p> <p>The value or object whose type is incorrect in this context. (The first version is used when the value in error was passed as an Any, and the second when it was passed as a RefObject.)</p> <p>“expected_type” : Any(TypeCode) The CORBA TypeCode that the value should have been.</p>	<p>“Wrong Type” arises when a RefObject or an Any value has the wrong type for context in which it was supplied. For example;</p> <ul style="list-style-type: none"> • A RefObject whose most derived type is incorrect; e.g. has the wrong M2-level Class or is a Class proxy instead of Instance, or vice versa. • An Any value that contains a single value where a sequence is required, or vice versa. • An Any value that contains a single value or sequence of values of the wrong CORBA type.
Wrong Number Parameters	Class or Operation for which the wrong number of actual parameters was supplied.	<p>“number_expected” : Any(Unsigned Long) The expected number of actual parameters.</p>	<p>“Wrong Number Parameters” arises when a client calls “refCreateInstance” or “refInvokeOperation” with too few or too many parameters.</p>
Invalid Deletion	A nil object reference	none	<p>“Invalid Deletion” arises when a client calls “refDelete” on a meta-object that cannot be deleted this way; i.e. an Association object, a Class Proxy object or a dependent Package object.</p>

The following IDL defines the **error_kind** strings for the above Reflective errors:

```

const string INVALID_DESIGNATOR_VIOLATION =
    "org.omg.mof:reflective.invalid_designator";
const string WRONG_DESIGNATOR_DESIGNATOR_VIOLATION =
    "org.omg.mof:reflective.wrong_designator_kind";
const string UNKNOWN_DESIGNATOR_VIOLATION =
    "org.omg.mof:reflective.unknown_designator";
const string ABSTRACT_CLASS_VIOLATION =
    "org.omg.mof:reflective.abstract_class";
const string NOT_CHANGEABLE_VIOLATION =
    "org.omg.mof:reflective.not_changeable";
const string NOT_NAVIGABLE_VIOLATION =
    "org.omg.mof:reflective.not_navigable";
const string NOT_PUBLIC_VIOLATION =
    "org.omg.mof:reflective.not_public";
const string WRONG_SCOPE_VIOLATION =

```

```

"org.omg.mof:reflective.wrong_scope";
const string WRONG_MULTIPLICITY_VIOLATION =
"org.omg.mof:reflective.wrong_multiplicity";
const string WRONG_TYPE_VIOLATION =
"org.omg.mof:reflective.wrong_type";
const string WRONG_NUMBER_PARAMETERS_VIOLATION =
"org.omg.mof:reflective.wrong_number_parameters";
const string INVALID_DELETION_VIOLATION =
"org.omg.mof:reflective.invalid_deletion";

```

Other Exception

There is one exception to this. When an Operation defined in the metamodel raises an Exception that is also defined in the metamodel; see below.

The `OtherException` exception is raised when a call to “`refInvokeOperation`” results in an error condition that correspond to an M2-level Exception defined for the Operation in the metamodel.

```

exception OtherException {
    DesignatorType exception_designator;
    ValueType exception_args;
};

```

The arguments to the `OtherException` exception are as follows:

- **exception_designator** gives the designator for the M2-level Exception raised.
- **exception_args** is an ordered list of CORBA Any values that represent the arguments of the Exception raised. The encoding of this field is defined in the specification of ““The “`refInvokeOperation`” operation invokes a metamodel defined Operation on the Instance or Class proxy object with the arguments supplied.” on page 6-20.

Note – When an error condition could be expressed as either a Reflective error a Structural error, the latter takes precedence. For example, if one end of Link in a call to “`refAddLink`” is a nil object reference, this should be signalled as “Nil Object” rather than “Wrong Type”

5.5 *Preconditions for IDL Generation*

The IDL mapping may not produce valid CORBA IDL if any of the following preconditions on the input meta-model is not satisfied:

- The MOF Model constraints, as defined above, must all be satisfied for the input meta-model.
- The input meta-model must be structurally consistent.

- The visible Names within a NameSpace must conform to the standard CORBA IDL identifier syntax. The original Names of Model Elements that have a valid substitute Name are excepted from this precondition; (see Section 5.6.2.1, “Substitute Name,” on page 5-37). No such requirement applies to Model Elements such as Tags, TypeAliases and some DataTypes whose names are not visible in the IDL mapping. However, for these “invisible” elements it is advisable to use a naming convention that minimizes the risk of name collision.
- The visible Names within a NameSpace must be unique after name substitution (see Section 5.6.2.1, “Substitute Name,” on page 5-37) application of the Format1 or Format2 name rewriting algorithms and other name mangling specified in the mapping.
- A DataType’s “typeCode” must not be an anonymous non-primitive type. It cannot have a kind of tk_array, tk_sequence, tk_string, tk_wstring or tk_fixed (with the exception of the TypeCodes for the IDL “string” and “wstring” types).
- A DataType’s “typeCode” must follow the conventions for expressing types and linking them to their definitions that are described in Section 3.4.7, “DataType,” on page 3-33 and Section 3.4.8, “TypeAlias,” on page 3-35. In addition:
 - When DataType’s “name” does not start with “*”, its Format 1 rendering must be identical to the type name in the DataType’s “typeCode.”
 - Whenever some embedded type in a DataType’s “typeCode” is linked (via a TypeAlias) to a defining Classifier, the Format 1 rendering of the Classifier’s “name” must be identical to the embedded type’s name in the TypeCode.
 - If the Classifier in the previous case is a DataType, the Classifier’s “typeCode” value must be identical to the embedded type’s TypeCode.
- A nested Package may not be used as a subtype or supertype.
- A nested Package may not import or be imported by another Package.
- The following interim visibility definitions and constraints apply to the IDL mapping:
 - A ModelElement is visible to another ModelElement only if the former has visibility of “public_vis.”
 - A ModelElement declared within another top-level Package is visible within a top-level Package only if the former Package is imported, clustered or inherited by the latter Package.
 - One ModelElement can only depend on another (in the sense of the M2-level DependsOn Association) if the latter is visible from the former within the definition of visibility immediately above.
 - After name substitution (see Section 5.6.2.1, “Substitute Name,” on page 5-37), the name of an Import must equal the name of its “importedNamespace.”
- A Class may not be nested within another Class.
- A Class may not be imported.
- If a Constraint is contained by a DataType or Operation, its name must also be unique in the DataType or Operation’s container Namespace.
- Model Elements in a meta-model cannot be cyclically dependent except as follows:

- A dependency cycle consisting of one or more Classes is legal, provided they all have the same container.
- A dependency cycle consisting of one or more Classes and one or more DataTypes or Exceptions, is legal provided they all have the same container.

Note – This precludes circular importing and circular clustering. It also precludes recursion between “pure” DataTypes. (The two exceptions correspond to cases that can be expressed in OMG IDL using forward interface declarations.)

CORBA 2.3 adds an additional IDL constraint: “The name of an interface or a module may not be redefined within the immediate scope of the interface or the module.” For example:

```

module M {
    typedef short M; // Error: M is the name of the module
                                // in the scope of which the typedef is
    interface I {
        void i (in short j);
                                // Error: i clashes with the interface
    };
};

```

The IDL templates in this specification do not contain any patterns of this form. However, poor choice of names in a meta-model may generate IDL that violates this constraint. In particular, the same name should not be used for both a container and its contents. For example, a Package should not have the same name as one of its Classes, DataTypes, or Associations. A Class should not have the same name as one of its Attributes or References. An Association should not have the same name as one of its AssociationEnds.

5.6 Standard Tags for the IDL Mapping

This section defines the standard Tags that apply to the Model to IDL mapping. Other Tags may be attached to the elements of a meta-model, but the meaning of these Tags is not specified. Similarly, this section does not specify the meaning of the Tags below in contexts apart from the Model to IDL mapping.

All standard Tag identifiers for the IDL mapping start with the prefix string:

“org.omg.mof.idl_”

Table 5-6 shows the conventions used to describe the standard Tags and their properties.

Table 5-6 Notation for Describing Standard Tags

<i>tag id:</i>	A string that denotes the semantic category for the tag.
<i>attaches to:</i>	Gives the kind(s) of Model::ModelElement that this category of tag can be meaningfully attached to.
<i>values:</i>	Gives the number and types of the tag’s values (i.e., parameters), if any. (Tag parameters are expressed as an unordered sequence of CORBA “any” values.)
<i>meaning:</i>	Describes the meaning of the tag in this context.
<i>idl generation:</i>	Defines the tag’s impact on the generated IDL.
<i>restrictions:</i>	Tag usage restrictions - for example: “at most one tag of this kind per element,” or “tag must be contained by the meta-model.”

Note – Many of these Tags significantly alter the interface signatures of the generated IDL. It is prudent for an IDL generator to only respect IDL mapping Tags when they are contained within the respective meta-model. Otherwise, it may not be possible to determine which Tags were in effect when the meta-data server was generated. This would make it hard for a client to infer the meaning of generated IDL at runtime. It would also make problems for automatic server and client generators.

5.6.1 Tags for Specifying IDL #pragma prefix

This tag allows the meta-modeler to specify the CORBA Interface Repository Identifier prefix for the generated IDL. This is essential when a MOF meta-model is used as the authoritative source for IDL for some other OMG standard.

5.6.1.1 IDL Prefix

<i>tag id:</i>	“org.omg.mof.idl_prefix”
<i>attaches to:</i>	Model::Package
<i>values:</i>	a String
<i>meaning:</i>	This tag supplies a RepositoryId prefix that is used for the entire module generated for the Package.

<i>tag id:</i>	“org.omg.mof.idl_prefix”
<i>idl generation:</i>	A #pragma prefix is inserted into the IDL before the “module” declaration for the Package.
<i>restrictions:</i>	[1] A Prefix tag should only be attached to a non-nested Package. [2] A Prefix tag contained by a Package takes precedence over one that is not contained.

5.6.2 Tags for Providing Substitute Identifiers

There are some situations when the IDL identifiers produced by the IDL mapping templates will result in name collisions. The following tag allows a meta-modeler to provide a substitute for a model element’s name that will be used in IDL generation.

5.6.2.1 Substitute Name

<i>tag id:</i>	“org.omg.mof.idl_substitute_name”
<i>attaches to:</i>	Model::ModelElement
<i>values:</i>	a String
<i>meaning:</i>	The String is the substitute name to be used in place of the model element’s name.
<i>idl generation:</i>	Wherever the IDL mapping makes use of a model element’s name, the substitute name should be used in its place. This substitution occurs before application of format1, format2, and other name mangling.
<i>restrictions:</i>	The preconditions described in Section 5.5, “Preconditions for IDL Generation,” on page 5-33 apply to the substitute name. For example: [1] it must be a syntactically valid IDL identifier, and [2] all identifiers produced from it must be unique in their respective scopes after formatting and name mangling, as per the IDL mapping specification. [3] in addition, there should be at most one Substitute Name tag per ModelElement.

5.6.3 Tags for Specifying IDL Inheritance

The following tags allow the meta-modeler to specify that a generated interface inherits from one or more additional IDL interfaces. These tags allow the definition of MOF-based meta-models that are upwards compatible with pre-existing meta-data interfaces expressed in CORBA IDL.

5.6.3.1 Instance Supertypes

<i>tag id:</i>	“org.omg.mof.idl_instance_supertypes”
<i>attaches to:</i>	Model::Class
<i>values:</i>	one or more TypeCodes (order is significant)
<i>meaning:</i>	The TypeCodes identify one or more IDL interfaces that the “instance” interface for this Class should inherit from.
<i>idl generation:</i>	The specified interfaces are added to the “instance” interface’s inheritance list following the other supertypes defined by the templates. They appear in the order given.
<i>restrictions:</i>	[1] The TypeCodes must have kind of ‘tk_objref’. [2] There should be at most one Instance Supertypes tag per Class.

5.6.3.2 Class Proxy Supertypes

<i>tag id:</i>	“org.omg.mof.idl_class_proxy_supertypes”
<i>attaches to:</i>	Model::Class
<i>values:</i>	one or more TypeCodes (order is significant)
<i>meaning:</i>	The TypeCodes identify one or more IDL interfaces that the “class proxy” interface for this Class should inherit from.
<i>idl generation:</i>	The specified interfaces are added to the “class proxy” interface’s inheritance list following the other supertypes defined by the templates. They appear in the order given.
<i>restrictions:</i>	[1] The TypeCodes must have kind of ‘tk_objref’. [2] There should be at most one Class Proxy Supertypes tag per Class.

5.6.3.3 Association Supertypes

<i>tag id:</i>	“org.omg.mof.idl_association_supertypes”
<i>attaches to:</i>	Model::Association
<i>values:</i>	one or more TypeCodes (order is significant)

<i>meaning:</i>	The TypeCodes identify one or more IDL interfaces that the interface for this Association should inherit from.
<i>idl generation:</i>	The specified interfaces are added to the “association” interface’s inheritance list following the other supertypes defined by the templates. They appear in the order given.
<i>restrictions:</i>	[1] The TypeCodes must have kind of ‘tk_objref.’ [2] There should be at most one Association Supertypes tag per Association.

5.6.3.4 Package Supertypes

<i>tag id:</i>	“org.omg.mof.idl_package_supertypes”
<i>attaches to:</i>	Model::Package
<i>values:</i>	one or more TypeCodes (order is significant)
<i>meaning:</i>	The TypeCodes identify one or more IDL interfaces that the interface for this Package should inherit from.
<i>idl generation:</i>	The specified interfaces are added to the “package” interface’s inheritance list following the other supertypes defined by the templates. They appear in the order given.
<i>restrictions:</i>	[1] The TypeCodes must have kind of ‘tk_objref.’ [2] There should be at most one Package Supertypes tag per Package.

5.7 Generated IDL Issues

During the design of the MOF Model to IDL mapping, several design decisions were made which are explained in this section.

5.7.1 Generated IDL Identifiers

Identifier naming is an important issue for automatically generated IDL, especially when that IDL is intended to be used by applications written by human programmers. The mapping has to reach a balance between conflicting requirements:

- Syntactic correctness - all identifiers in the mapped IDL must conform to the defined CORBA IDL syntax, and they must all conform to the CORBA scoping and upper/lower casing restrictions.
- User friendliness - identifiers should convey as much information as possible without being overly long.
- Conformance to existing conventions - identifiers should conform to existing stylistic conventions.

The OMG conventions for IDL identifiers (see “OMG IDL Style Guide: ab/98-06-03”) are based on the notion that an identifier is formed from one or more words in some natural language. The conventions allow digits to be used in words and take account of acronyms. The Style Guide then specifies three different styles for putting some words together as an identifier. In particular:

- Identifiers for IDL module, interface and types are capitalized. If the name consists of multiple words, each word is capitalized in the identifier.
- Identifiers for IDL operations, attributes, formal parameters, struct and exception members, and union branches are all lower-case. If the name consists of multiple words, the words are separated by underscores (“_”) in the identifier.
- Identifiers for IDL constant and enumerator names are all upper-case. If the name consists of multiple words, the words are separated by underscores (“_”) in the identifier.

5.7.1.1 *Rules for splitting MOF Model::ModelElement names into "words"*

According to the MOF Model, the “name” of a ModelElement is an instance of the NameType (i.e., a CORBA string). With a small number of exceptions, the IDL mapping needs to be able to convert these NameType instances into CORBA IDL identifiers for use in a variety of contexts.

Since the MOF Model (like the UML meta-model) does not restrict the strings that can be used as ModelElement “name,” mapping them to meaningful IDL identifiers is not possible in the most general case. For example, names that include graphic characters or accented letters do not map to IDL identifiers.

The IDL mapping requires that those names needing to be mapped consist only of unaccented upper- and lower-case Latin letters, the digits ‘0’ to ‘9,’ hyphens (‘-’), underscores (‘_’), and white-space characters. In addition, it must be possible to split a name into “words” according to the following specification.

A “word” is defined to be an upper-case letter, followed optionally by more upper-case letters and digits and then optionally by lower-case letters and digits. An underscore (“_”), hyphen (“-”) or white space character will terminate a word. This is expressed more formally by the following mini-grammar:

```

word      ::=    [A-Z][A-Z0-9]*[a-z0-9]*
              |    [a-z][a-z0-9]*
whitespace ::=    SP, CR, NL, HT, VT, etc
term      ::=    { '_' | '-' | whitespace }*

identifier ::=    [term] word { [term] word }* [term]

```

The sequence of “words” for a name can then be formed into OMG IDL identifiers according to the 3 formats below.

5.7.1.2 IDL Identifier Format 1

In Format 1, the first letter of each word is converted into upper case, and other letters remain the same case as input. The words are not separated by other characters. Table 5-7 lists some examples of Format 1 identifiers.

Table 5-7 Format 1 Identifiers

Name	Name split into words	Identifier in Format 1
foo	"foo"	Foo
foo_bar	"foo" "bar"	FooBar
ALPHAbeticalOrder	"ALPHAbetical" "Order"	ALPHAbeticalOrder
-a1B2c3-d4-	"a1" "B2c3" "d4"	A1B2c3D4
DSTC pty ltd	"DSTC" "pty" "ltd"	DSTCPtyLtd

Format 1 is used by the IDL mapping to produce the names of modules and interfaces.

5.7.1.3 IDL Identifier Format 2

In Format 2, all letters in each word are converted into lower case. Each word is separated by an underscore "_". Table 5-8 lists some examples of Format 2 identifiers.

Table 5-8 Format 2 Identifiers

Name	Name split into words	Identifier in Format 2
foo	"foo"	foo
foo_bar	"foo" "bar"	foo_bar
ALPHAbeticalOrder	"ALPHAbetical" "Order"	alphabetical_order
-a1B2c3_d4_	"a1" "B2c3" "d4"	a1_b2c3_d4
DSTC pty ltd	"DSTC" "pty" "ltd"	dstc_pty_ltd

Format 2 is used by the IDL mapping for identifiers for IDL operations, exceptions, attributes, formal parameters, exception members and members of generated struct types.

5.7.1.4 IDL Identifier Format 3

In Format 3, all letters in each word are converted into upper case. Each word is separated by an underscore "_". Table 5-9 lists some examples of Format 3 identifiers.

Table 5-9 Format 3 Identifiers

Name	Name split into words	Identifier in Format 3
foo	"foo"	FOO
foo_bar	"foo" "bar"	FOO_BAR
ALPHAbeticalOrder	"ALPHAbetical" "Order"	ALPHABETICAL_ORDE R
-a1B2c3_d4_	"a1" "B2c3" "d4"	A1_B2C3_D4
DSTC pty ltd	"DSTC" "pty" "ltd"	DSTC_PTY_LTD

Format 3 is used by the IDL mapping for identifiers for IDL constants.

5.7.1.5 Identifiers in TypeCodes

DataTypes are used to specify various kinds of types for use in Attribute and Parameter definitions. These types are encoded using the CORBA TypeCode type, and contain embedded names of types, struct members, union arms and enumerators.

Unlike ModelElement names, the names embedded in TypeCodes are constrained by the CORBA Core specification to be valid identifiers. For this reason alone, they are not reformatted by the IDL mapping. (If the IDL mapping did reformat these names, the resulting IDL would declare types that do not match the TypeCode. Among other things, a reflective client would need to reinterpret the DataType's "typeCode" field in order to produce an Any value of the correct type. This is clearly undesirable.)

There are a couple of cases that require further specification:

- When the IDL mapping produces a qualified name for an external type encoded as a DataType, the components of the name are not subject to reformatting.
- When the IDL mapping produces collections types corresponding to a DataType, the corresponding identifiers are formatted according to the Format 1 rules.

5.7.1.6 Literal String Values

Literal string values (in string valued Constants) are not re-formatted and appear in the generated IDL exactly as specified by the Constant's "value" attribute.

5.7.2 Generation Rules for Collection Types

The MOF Model allows Attributes, AssociationEnds, References and Parameters to being single-, optional- or multi-valued depending on the ModelElement's base type and its multiplicity.

At various places in the mapped interfaces, it is necessary to pass collections that represent values for the optional- or multi-valued cases. The IDL types for such a collection is a **typedef** alias for an unbounded CORBA sequence of the collection base type. The name of the typedef depends on the corresponding ModelElement's multiplicity specification.

For example, if the ModelElement is ordered and unique, then the collection type is a unique list (or UList). The typedef name for a unique list takes the form `<ClassifierType>UList` (i.e., the name of the collection base type followed by the characters "UList"). For example, if an M2-level Operation returns an ordered, unique list of Class "Foo," then IDL result type for the corresponding operation to be called "FooUList" with the declaration:

```
typedef sequence <Foo> FooUList;
```

There are four distinct collection type suffixes corresponding to the combinations of the "isOrdered" and "isUnique" flags for an element's "multiplicity" attribute. The appropriate suffix should be generated whenever `<CollectionKind>` appears in the IDL templates below.

Table 5-10 Collection Kinds

Multiplicity Flags	Collection	Kind Suffix
none	bag	Bag
ordered	list	List
unique	set	Set
ordered, unique	unique list (ordered set)	UList

Note that the MOF Model specification includes a relevant Constraint on multiplicity values; see the "MustBeUnorderedNonunique" constraint in Section 3.9.4, "The MOF Model Constraints," on page 3-91. This states that when a feature's multiplicity bounds are [0..1], both the "isOrdered" and "isUnique" are set to false. As a consequence, the `<CollectionKind>` suffix for a [0..1] collection type is always "Bag".

Similar collection kind naming conventions are used for DataTypes. Thus for a set of some enumeration type, the mapping would produce the following:

```
enum SomeEnum {e1, e2};  
typedef sequence <SomeEnum> SomeEnumSet;
```

When the DataType is a built-in CORBA type, the base name for the type is defined as follows:

Table 5-11 Base Names for Built-in CORBA Types

Built-in CORBA type	Base name
short	Short
long	Long

Table 5-11 Base Names for Built-in CORBA Types

unsigned short	UShort
unsigned long	ULong
float	Float
double	Double
boolean	Boolean
char	Char
string	String
octet	Octet
any	Any
TypeCode	TypeCode
Object	Object

The declarations for collection types will appear in one of three places.

- If the collection's base type is defined somewhere within the top-level Package being generated, collection type declarations appear immediately following the base type's introduction; see Section 5.8.16, "DataType Template," on page 5-90 or Section 5.8.5, "Class Forward Declaration Template," on page 5-53.
- If the base type is imported or inherited from another Package, the collection type declarations appear at the beginning of the IDL module for the Package that imports or inherits the type's Package; see Section 5.8.2, "Package Module Template," on page 5-46.
- If the base type is a CORBA built-in type, or a CORBA type with an external (non-MOF) declaration, the collection type declarations appear at the beginning of the IDL module for the outermost Package; see Section 5.8.2, "Package Module Template," on page 5-46.

Since CORBA sequence types require considerable run-time support code in some language bindings, collection type declarations must only be generated if they are needed within the IDL for the current outermost Package.

Operations produced by the IDL mapping with collection parameters must ensure that the sequence values supplied and returned have an appropriate number of elements. When collection parameters are sets or unique lists, operations must also ensure that the sequence values contains no duplicates.

5.7.3 IDL Identifier Qualification

To avoid scoping errors within the mapped IDL, identifier names must be either fully qualified, or partially qualified to an appropriate level. This specification leaves the choice between the use of fully or partially qualified identifiers to the implementer.

5.7.4 File Organization and #include statements

This specification does not prescribe how the generated IDL is organized into files. Therefore, the generation rules do not contain any “#include” statements. An implementer must decide how to organize the generated IDL into files, and must generate appropriate “#include” statements to ensure that the resultant IDL can compile. Similarly, the implementer must generate “#ifndef” guards as required by the OMG style rules.

5.8 IDL Mapping Templates

Model specific IDL is produced by traversing the containment hierarchy of a top-level M2-level Package. The CORBA module structure of the resulting IDL directly reflects the containment hierarchy of the source Package. If element X contains element Y in the source model, then the IDL corresponding to X will have the IDL corresponding to Y embedded in it (assuming that IDL is produced for Y).

The IDL mapping supports the containment hierarchy for ModelElements as described in Section 3.3.4, “The MOF Model Containment Hierarchy,” on page 3-14, except as stated in Section 5.5, “Preconditions for IDL Generation,” on page 5-33. Further restrictions on meta-models that can be successfully mapped are described in the same section.

The mapping rules are described in terms of IDL templates. Each Template describes the maximum IDL which could be generated when mapping MOF Model objects. In any specific case, the actual IDL generated will depend on the properties of the corresponding MOF Model object.

Throughout the following Template descriptions, the IDL is said to be “generated by” the Templates. Clearly the Templates do not generate IDL in a literal sense. Instead, the reader should imagine that each Template is a parameter to a hypothetical generator function. When it is called with the appropriate kind of MOF ModelElement object as a second parameter, the function “elaborates” the template to produce an appropriate fragment of CORBA IDL. A similar “elaboration” process gives the required semantics for the IDL from the descriptions following the templates and the specifications given earlier in Section 5.3, “Computational Semantics for the IDL Mapping,” on page 5-6.

Note – The Template approach used here is a notational convenience, not a required or suggested implementation strategy.

5.8.1 Template Notation

The following table is a guide to interpreting the IDL generation templates.

Table 5-12 IDL Generation Templates Guide

Appearance (by example)	Meaning
typedef	The literal characters in bold font should be generated.
<AttributeType>	The characters should be substituted for the described identifier using Identifier Format 1. The <> do not appear in the generated IDL.
<attribute_name>	The characters should be substituted for the described identifier using the Identifier Format 2. The <> do not appear in the generated IDL.
<CONSTANT_NAME>	The characters should be substituted for the described identifier using the Identifier Format 3. The <> do not appear in the generated IDL.
<CONSTANTVALUE>	The characters should be substituted for the described identifier without formatting (i.e., as is). Typically, these are literal values. The <> do not appear in the generated IDL.
<<XYZ TEMPLATE>>	Apply the named template. The <<>> do not appear in the generated IDL.
<i>some phrase , . . .</i>	The ellipsis characters “. . .” following the “,” indicate that this generates a comma separated list of “some phrase”. It is implicit that there is no comma at the end of the list.
[<i>some phrase</i>]	The square bracket characters “[]” surrounding a phrase in a template indicate that the phrase may or may not be required, depending on context.
<i>// for each parameter</i>	Gives the rules on when and how to perform the IDL generation, or some general commentary on the process. The rules themselves do not appear in the generated IDL.

5.8.2 Package Module Template

This section describes the rules for mapping an MOF Package object to a CORBA IDL module as expressed in the Package Module Template.

The Package Module Template generates a CORBA IDL module that contains the IDL for each of the M2-level Constants, DataTypes, Exceptions, Constraints, Imports, Classes, and Associations in an M2-level Package. It also contains the IDL for the M1-

level Package and Package Factory interfaces, and type declarations for various collection types. Most of this is defined in subsidiary templates. IDL generation is suppressed if the Package “visibility” is not “public_vis.”

5.8.2.1 *emplate*

```

<<ANNOTATION TEMPLATE>>
// if this Package has visibility of private or protected, no IDL is
// generated for it
module <PackageName> {

// if this Package is a "top-level" Package, generate any collection types
// for built-in types that are required by the IDL for this Package and
// its contents
typedef sequence < <BuiltinType> > <BuiltinTypeName><CollectionKind>;

// if the Package has superPackages, generate any collection
// types for inherited types that are required by the IDL for this Package
// and its contents
typedef sequence < <InheritedType> > <InheritedType><CollectionKind>;

// if the Package has Imports, generate any collection types for
// imported or clustered types that are required by the IDL for this Package
// and its contents
typedef sequence < <ImportedType> > <ImportedType><CollectionKind>;

interface <PackageName>Package; // forward declaration

// for each Class contained in the Package
<<CLASS FORWARD DECLARATION TEMPLATE>>

// for each Package, DataType, Exception, Class, Association, Constraint,
// and Constant contained by the Package, generate the appropriate IDL
<<PACKAGE MODULE TEMPLATE>>
<<DATATYPE TEMPLATE>>
<<EXCEPTION TEMPLATE>>
<<CLASS TEMPLATE>>
<<ASSOCIATION TEMPLATE>>
<<CONSTRAINT TEMPLATE>>
<<CONSTANT TEMPLATE>>

// Generate the Package Factory interface

```

```
<<PACKAGE FACTORY TEMPLATE>>

// Generate the Package interface
<<PACKAGE TEMPLATE>>

}; // end of module <PackageName>
```

Description

The Package Module Template starts by rendering the M2-level Package's "annotation" attribute as a comment using the Annotation Template. This is followed by the IDL module header for the Package's module. The module name is *<PackageName>*.

The first group of declarations within the module are the sequence type declarations for collection types:

- If the M2-level Package is a top-level Package, the template generates collection types for the MOF built-in types as required.
- If the M2-level Package has superPackages, the template generates collection types for inherited Classes and DataTypes as required.
- If the M2-level Package has Imports, the template generates collection types for any Classes and DataTypes in the imported or clustered Packages.

Note – Collection types should only be generated if they are going to be used within the current outermost module.

After the collection types, the template generates forward declarations for some IDL interfaces. First, it forward declares the M1-level Package interface, giving it the name *<PackageName>Package*. Then, it forward declares the Class proxy and Instance interfaces for all M2-level Classes in the current M2-level Package's "contents" using the template defined in Section 5.8.5, "Class Forward Declaration Template," on page 5-53.

Next, IDL must be generated for the current M2-level Package's "contents" as follows:

- For nested Packages, use the template defined in Section 5.8.2, "Package Module Template," on page 5-46.
- For Classes, use the template defined in Section 5.8.6, "Class Template," on page 5-54.
- For Associations, use the template defined in Section 5.8.10, "Association Template," on page 5-59.
- For Constants, use the template defined in Section 5.8.13, "Operation Template," on page 5-86.
- For Exceptions, use the template defined in Section 5.8.14, "Exception Template," on page 5-89.

- For DataTypes, use the template defined in Section 5.8.16, “DataType Template,” on page 5-90.
- For Constraints, use the template defined in Section 5.8.17, “Constraint Template,” on page 5-91.

The IDL for the contained ModelElements must be generated in an order that reflects their dependencies. For example, the IDL for a DataType should appear before the IDL for other ModelElements that use it.

Finally, the Package Module Template generates the Package Factory and Package interfaces for the current M2-level Package using the templates respectively defined in Section 5.8.3, “Package Factory Template,” on page 5-49 and Section 5.8.4, “Package Template,” on page 5-51.

5.8.3 Package Factory Template

The Package Factory Template defines the IDL generation rules for the Package Factory interface; see Section 5.2.1.1, “Package objects and Package Factory objects,” on page 5-2 and Section 5.2.2, “The Meta Object Interface Hierarchy,” on page 5-4.

A Package Factory interface is generated for top-level M2 Packages only. The interface is named `<PackageName>PackageFactory` and it contains a single “factory” operation, as described below.

Template

```
// if the this Package is top-level
interface <PackageName>PackageFactory
{
    <PackageName>Package create_<package_name>_package (
        // for each non-derived class-level Attribute of any directly or
        // indirectly contained Class within this Package and its closure
        // under Package generalization and clustering.
        in <AttributeType>[<CollectionKind>]
            <qualified_attribute_name>, ...
    )
    raises (Reflective::MofError);
};
```

IDL Supertypes

none

Operations

create_<package_name>_package

The “create_<package_name>_package” operation creates a new Package object that is an instance of this M2-level Package.	
<i>reflective analog:</i>	none
<i>return type:</i>	<PackageName>Package
<i>parameters:</i>	<qualified_attribute_name> : in <AttributeType>[<CollectionKind>], ...
<i>exceptions:</i>	MofError (Overflow, Underflow, Duplicate)

The parameters for “create_<package_name>_package” give the initial values for any non-derived classifier-scoped Attributes for all Classes that belong to this M2-level Package’s extent.

As Attributes in different Classes can have the same name, the parameter name <qualified_attribute_name> is qualified relative to the Package (e.g., “class1_attribute1”).

When the Attribute multiplicity is not [1..1], the <AttributeType> has an appropriate CollectionKind suffix appended; see Section 5.7.1.6, “Literal String Values,” on page 5-42.

The parameters are declared in a sequence defined by a recursive depth-first traversal of the Package's ancestors clusters and components, visiting a Package's supertypes before its contents. The following ordering rules apply:

1. A Package’s supertype Packages are processed before the “contents” of the Package.
2. The supertype Packages are processed in the order defined by the “Generalizes” association.
3. Classes, Imports (with “isClustered” set to true) and nested Packages within a Package are processed in the order of the “Contains” association.
4. A Class’s superclasses are processed before the “contents” of the Class.
5. Any Class superclasses are processed in the order of the “Generalizes” association.
6. An Import with “isClustered” set to true is processed by processing the clustered Package.
7. Attributes within a Class are processed in the order of the “Contains” association.
8. When an Attribute is encountered that has already been encountered during the traversal, generation of another initialization parameter is suppressed.

The MofError exception can be raised if there is an Structural, Constraint or Semantic errors. In particular, “Overflow,” “Underflow,” and “Duplicate” occur if an Attribute initialization parameter does not conform to the respective Attribute’s multiplicity specification.

5.8.4 Package Template

The Package Factory Template defines the IDL generation rules for the Package interface; see Section 5.2.1.1, “Package objects and Package Factory objects,” on page 5-2 and Section 5.2.2, “The Meta Object Interface Hierarchy,” on page 5-4.

A Package interface is named `<PackageName>Package` and it contains read-only IDL attributes giving the dependent Package, Association and Class proxy objects for a Package object.

Template

```

interface <PackageName>Package :
    // if Package has no super-Packages
    Reflective::RefPackage
    // else for each public super-Package (in order)
    <SuperPackage>Package, ...
    // if Package has a "Package Supertypes" Tag
    //     for each supertype defined by the Tag (in order)
    , <PackageSupertypeName>, ...
{
    // for each Package for an Import where:
    //     is_clustered == true and
    //     Import.visibility == public and
    //     importedNamespace.visibility == public
    readonly attribute <ClusteredPackageName>Package
        <clustered_package_name>_ref;
    // for each public contained Package
    readonly attribute <NestedPackageName>Package
        <nested_package_name>_ref;
    // for each public contained Class
    readonly attribute <ClassName>Class <class_name>_ref;
    // for each public contained Association
    readonly attribute <AssociationName> <association_name>_ref;
};

```

Supertypes

If the M2-level Package inherits from other M2-level Packages with “visibility” of “public_vis,” the Package interface inherits from the interfaces corresponding super-Packages. Otherwise, the Package interface inherits from **Reflective::RefPackage**.

If the M2-level Package has a “Package Supertypes” Tag (see Section 5.6.3, “Tags for Specifying IDL Inheritance,” on page 5-37), the generated Package interface also inherits from the IDL interfaces specified by the Tag.

Attributes

clustered_package_name>_ref

An attribute of this form is generated for each public clustered Package of the current M2-level Package. The attribute is generated if and only if

1. the Import’s “isClustered” flag is true,
2. the Import’s “visibility” is “public_vis”,
3. the Import’s “importedNamespace” is a Package, and
4. the clustered Package has “visibility” of “public_vis.”

The attribute holds the object reference for the M1-level Package’s M1-level clustered Package object.

<i>reflective analog:</i>	ref_package_ref(<clustered_package_designator>); Section , “Reflective::RefPackage,” on page 6-28.
<i>type:</i>	<ClusteredPackageName> Package
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	no

<nested_package_name>_ref

An attribute of this form is generated for each nested Package in the current M2-level Package whose “visibility” is “public_vis”. The attribute holds the object reference for the M1-level Package’s M1-level nested Package object.

<i>reflective analog:</i>	ref_package_ref(<nested_package_designator>); Section , “Reflective::RefPackage,” on page 6-28.
<i>type:</i>	<NestedPackageName> Package
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	no

<class_name>_ref

An attribute of this form is generated for each Class in the current Package whose “visibility” is “public_vis.” The attribute holds the object reference for the M1-level Package’s M1-level Class Proxy object.

<i>reflective analog:</i>	ref_class_ref(<class_designator>)
<i>type:</i>	<ClassName>Class
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	no

<association_name>_ref

An attribute of this form is generated for each Association in the current Package whose “visibility” is “public_vis.” The attribute holds the object reference for the M1-level Package’s M1-level Association object.

<i>reflective analog:</i>	ref_package_ref(<association_designator>);
<i>type:</i>	<AssociationName>
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	no

Operations

none

5.8.5 Class Forward Declaration Template

The Class Forward Declaration Template defines the IDL generation rules for the forward interface declarations for an M2-level Class whose “visibility” is “public_vis.” It also produces any Class collection type declarations required by the IDL of the containing Package(s).

Template

```
// if the Class has visibility of protected or private, no IDL
// is generated.
```

```
interface <ClassName>Class;
```

```
interface <ClassName>;
```

```
// generate type declarations for any collections of this Class that
// will be used by the IDL being generated
```

```
typedef sequence < <ClassName> > <ClassName>Set;
```

```
typedef sequence < <ClassName> > <ClassName>Bag;  
typedef sequence < <ClassName> > <ClassName>List;  
typedef sequence < <ClassName> > <ClassName>UList;
```

Description

The Class Forward Declaration Template generates a forward declaration for the Instance and Class proxy interfaces for an M2-level Class. These have IDL identifiers *<ClassName>* and *<ClassName>Class* respectively. If any collection types for this Class are required, their declarations should follow the forward declarations.

Note – Collection types should only be generated if they are going to be used within the current outermost module.

5.8.6 Class Template

The Class Template defines the IDL generation rules for an M2-level Class whose “visibility” is “public_vis.” The IDL is generated within the module for the Class’s containing Package and consists of a comment followed by the complete Classes Class Proxy and Instance interfaces.

Template

```
// if the Class has visibility of protected or private, no IDL  
// is generated
```

```
<<ANNOTATION TEMPLATE>>  
<<CLASS PROXY TEMPLATE>>  
<<INSTANCE TEMPLATE>>
```

Description

See Section 5.8.7, “Class Proxy Template,” on page 5-54 and Section 5.8.6, “Class Template,” on page 5-54.

5.8.7 Class Proxy Template

The Class Proxy Template defines the IDL generation rules for the *<ClassName>Class* interface for an M2-level Class whose “visibility” is “public_vis.” This interface has operations for any classifier-scoped Attributes and Operations, along with a factory operation and IDL attributes that give access to the extant Instance objects. It also contains IDL declarations corresponding to any DataTypes, Exceptions, Constants, and Constraints in the M2-level Class.

Template

```

interface <ClassName>Class :
    // if Class has no super-Classes
    Reflective::RefObject
    // else for each super-Class
    <SuperClass>Class, ...
    // if Class has a "Class Proxy Supertypes" Tag
    //     for each supertype defined by the Tag (in order)
    , <ClassProxySupertypeName>, ...
{
    // all <ClassName> including subclasses of <ClassName>
    readonly attribute <ClassName>Set all_of_type_<class_name>;

    // if the Class is not abstract
    // all <ClassName> excluding subclasses of <ClassName>
    readonly attribute <ClassName>Set all_of_class_<class_name>;

    // for each Constant, DataType, Exception, Constraint,
    // classifier-scoped Attribute and classifier-scoped Operation
    // in the Class, generate the appropriate IDL
    <<DATATYPE TEMPLATE>>
    <<CONSTRAINT TEMPLATE>>
    <<CONSTANT TEMPLATE>>
    <<EXCEPTION TEMPLATE>>
    <<ATTRIBUTE TEMPLATE>> // public classifier-level only
    <<OPERATION TEMPLATE>> // public classifier-level only

    // if the Class is not abstract
    <<CLASS CREATE TEMPLATE>>
}; // end of interface <ClassName>Class

```

Supertypes

If the M2-level Class inherits from other M2-level Classes, the generated Class Proxy interface inherits from the corresponding supertype Class Proxy interfaces. Otherwise, the Class Proxy interface inherits from **Reflective::RefObject**.

If the M2-level Class has a “Class Proxy Supertypes” Tag (see Section 5.6.3, “Tags for Specifying IDL Inheritance,” on page 5-37), the generated Class Proxy interface also inherits from the IDL interfaces specified by the Tag.

Attributes

all_of_class_<class_name>

The “all_of_class_<class_name>” attribute gives all Instance objects in the current extent for the corresponding M2-level Class. The attribute is only generated if “isAbstract” is false for the Class.

<i>reflective analog:</i>	ref_all_objects(<class_designator>, false)
<i>type:</i>	<ClassName> (multiplicity zero or more, unique, non ordered)
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	no

The value of this attribute mirrors the definition of Instance object lifetimes; see Section 5.3.2.2, “Instance object lifecycle semantics,” on page 5-9. It does not include any deleted Instance objects.

all_of_type_<class_name>

The “all_of_type_<class_name>” attribute gives all Instance objects in the current extent for the corresponding M2-level Class or for any M2-level subClasses.

<i>reflective analog:</i>	ref_all_objects(<class_designator>, true)
<i>type:</i>	<ClassName> (multiplicity zero or more, unique, non ordered)
<i>multiplicity:</i>	exactly one
<i>changeable:</i>	no

The value of this attribute mirrors the definition of Instance object lifetimes; see Section 5.3.2.2, “Instance object lifecycle semantics,” on page 5-9. It does not include any deleted Instance objects.

Operations

The operations for a <ClassName>Class interface are produced by the Attribute, Operation and Class Create Templates. Note that the operations for the M2-level Classes instance-scoped features do not appear in this interface.

5.8.8 Instance Template

The Instance Template defines the IDL generation rules for the <ClassName> interface for an M2-level Class whose “visibility” is “public_vis.” This interface contains operations for the M2-level Classes instance-scoped Attributes and Operations, along with any References.

Template

```

interface <ClassName> :
    // (The Instance interface inherits the Class Proxy interface
    // for the Class and Instance interfaces for any super-Classes)
    <ClassName>Class
    // for each super-Class of this Class (in order)
    , <SuperClassName>, ...
    // if Class has an "Instance Supertypes" Tag
    //     for each supertype defined by the Tag (in order)
    , <InstanceSupertypeName>, ...
{
    // for each Attribute, Reference, Operation contained in
    // this Class, generate the appropriate IDL
    <<ATTRIBUTE TEMPLATE>> // public instance-level only
    <<REFERENCE TEMPLATE>> // public only
    <<OPERATION TEMPLATE>> // public instance-level only
}; // end of interface <ClassName>

```

Supertypes

The Instance interface for an M2-level Class inherits from the Class's Class Proxy interface, along with the Instance interfaces for all of its M2-level super-Classes.

If the M2-level Class has an "Instance Supertypes" Tag (see Section 5.6.3, "Tags for Specifying IDL Inheritance," on page 5-37), the generated Instance interface also inherits from the IDL interfaces specified by the Tag.

Attributes

none

Operations

The operations for an Instance interface are generated by the Attribute, Reference, and Operation Templates. Note that the operations for instance-scoped Attributes and Operations only appear here.

5.8.9 Class Create Template

The Class Create Template defines the IDL generation rules for the Instance factory operation for a non-abstract M2-level Class whose "visibility" is "public_vis."

Template

```

<ClassName> create_<class_name> (
    // for each non-derived direct or inherited attribute
    in <AttributeType>[<CollectionKind>] <attribute_name>, ...
)
raises (Reflective::MofError);

```

Operations

create_<class_name>

The “create_<class_name>” operation creates new Instance objects for the M2-level Class (i.e., instances of the Class's <ClassName> interface).

<i>reflective analog:</i>	ref_create_instance(<class_designator>, <attr_name>,...)
<i>return type:</i>	<ClassName>
<i>parameters:</i>	in <AttrTypeName>[<CollectionType>] <attr_name>, ...
<i>exceptions:</i>	MofError (Overflow, Underflow, Duplicate, Composition Closure, Supertype Closure, Already Created)

The parameters to “create_<class_name>” provide initial values for the M2-level Class's non-derived attributes. Parameter declarations are generated in an order defined by a recursive depth-first traversal of the inheritance graph. More precisely,

1. a Classes' super-Classes are processed before the Classes' Attributes,
2. super-Classes are processed in the order of the “Generalizes” association,
3. the Attributes of each Class or super-Class are processed in the order of the “Contains” association,
4. when an Attribute is encountered with a “scope” value of “classifier_level” or an “isDerived” value of true no parameter is generated, and
5. when an Attribute is encountered a second or subsequent time, no additional parameter is generated.

When an Attribute has multiplicity bounds other than [1..1], the type of the corresponding initial value parameter's type will be a collection type; see Section 5.7.1.6, “Literal String Values,” on page 5-42.

“Overflow,” “Underflow,” and “Duplicate” occur if an argument that gives the initial value for an Attribute does not match the Attribute's multiplicity specification.

“Composition Closure” occurs if the initial value for a “composite” Attribute contains an Instance object in another extent; see Section 4.9.2, “The Composition Closure Rule,” on page 4-20.

“Supertype Closure” occurs if the extent for the current object cannot create Instance objects for this super-Class; see Section 5.3.8, “The Supertype Closure Rule,” on page 5-22.

“Already Created” occurs if the M2-level Class has “isSingleton” set to true, and this object’s extent already includes an Instance object for the Class.

5.8.10 Association Template

The Association Template defines the generation rules for the Association interface corresponding to an M2-level Association whose “visibility” is “public_vis.” This interface contains the IDL operations for accessing and updating the Association’s M1-level link set.

Template

```
// If the Association has visibility of protected or private,
// no IDL is generated

// data types for Association <AssociationName>
struct <AssociationName>Link {
    <AssociationEnd1ClassName> <associationend1_name>;
    <AssociationEnd2ClassName> <associationend2_name>;
};
typedef sequence < <AssociationName>Link >
    <AssociationName>LinkSet;

<<ANNOTATION TEMPLATE>>

interface <AssociationName> : Reflective::RefAssociation
    // if Association has an "Association Supertypes" Tag
    //     for each supertype defined by the Tag (in order)
    , <AssociationSupertypeName>, ...

{
    // list of associated elements
    <AssociationName>LinkSet all_<association_name>_links ()
    raises (Reflective::MofError);

boolean exists (
```

```

        in <AssociationEnd1Class> <association_end1_name>,
        in <AssociationEnd2Class> <association_end2_name>)
        raises (Reflective::MofError);

// if association_end1 is_navigable
<AssociationEnd1Class>[<CollectionKind>] <association_end1_name> (
    in <AssociationEnd2Class> <association_end2_name>)
    raises (Reflective::MofError);

// if association_end2 is_navigable
<AssociationEnd2Class>[<CollectionKind>] <association_end2_name> (
    in <AssociationEnd1Class> <association_end1_name>)
    raises (Reflective::MofError);

// if association_end1 is_changeable
// and association_end2 is_changeable
void add (
    in <AssociationEnd1Class> <association_end1_name>,
    in <AssociationEnd2Class> <association_end2_name>)
    raises (Reflective::MofError);

// if association_end1 is_changeable and is_navigable
// and association_end2 is_changeable
// and association_end1 has upper > 1 and is_ordered
void add_before_<association_end1_name> (
    in <AssociationEnd1Class> <association_end1_name>,
    in <AssociationEnd2Class> <association_end2_name>,
    in <AssociationEnd1Class> before)
    raises (Reflective::NotFound, Reflective::MofError);

// if association_end1 is_changeable
// and association_end2 is_changeable and is_navigable
// and association_end2 has upper > 1 and is_ordered
void add_before_<association_end2_name> (
    in <AssociationEnd1Class> <association_end1_name>,
    in <AssociationEnd2Class> <association_end2_name>,
    in <AssociationEnd2Class> before)
    raises (Reflective::NotFound, Reflective::MofError);

// if association_end1 is_navigable and is_changeable
void modify_<association_end1_name> (

```



```

        in <AssociationEnd1Class> <association_end1_name>,
        in <AssociationEnd2Class> <association_end2_name>,
        in <AssociationEnd1Class> new_<association_end1_name>)
        raises (Reflective::NotFound, Reflective::MofError);
// if association_end2 is_navigable and is_changeable
void modify_<association_end2_name> (
    in <AssociationEnd1Class> <association_end1_name>,
    in <AssociationEnd2Class> <association_end2_name>,
    in <AssociationEnd2Class> new_<association_end2_name>)
    raises (Reflective::NotFound, Reflective::MofError);

// if association_end1 is_changeable
// and association_end2 is_changeable
void remove (
    in <AssociationEnd1Class> <association_end1_name>,
    in <AssociationEnd2Class> <association_end2_name>)
    raises (Reflective::NotFound, Reflective::MofError);
};

```

DataTypes

The Association Template generates data type declarations to that are used to pass a link set for the M2-level Association. The <AssociationName>**Link** and <AssociationName>**LinkSet** type declarations precede the Association interface declaration.

Supertypes

Every generated Association interface inherits from **Reflective::RefAssociation**. If the M2-level Association has an “Association Supertypes” Tag (see Section 5.6.3, “Tags for Specifying IDL Inheritance,” on page 5-37), the generated Association interface also inherits from the IDL interfaces specified by the Tag.

Attributes

none

Operations

all_<association_name>_links

The “all_<association_name>_links” operation creates new Instance objects for the M2-level Class (i.e., instances of the Class's <ClassName> interface).

<i>reflective analog:</i>	ref_all_links()
<i>return type:</i>	<AssociationName> LinkSet
<i>parameters:</i>	none
<i>query:</i>	yes
<i>exceptions:</i>	MofError()

The “all_<association_name>_links” operation returns the current link set for this Association expressed using the <AssociationName>**LinkSet** type.

While the definitions in Section 4.7.2.1, “A Mathematical Model of Association State,” on page 4-15 state that an ordered Association implies a partial ordering over the LinkSet, the result of the “all_<association_name>_links” operation is defined to be a Set. A client should not draw any conclusions from the ordering of the returned links.

The operation’s signature raises **Reflective::MofError** to allow Constraint error and Semantic error conditions to be signalled.

exists

The “exists” operation queries whether a link currently exists between a given pair of Instance objects in the current M1-level Association extent.

<i>reflective analog:</i>	ref_link_exists(Link{ <assoc_end1_name>, <assoc_end2_name> })
<i>return type:</i>	boolean
<i>parameters:</i>	in <AssocEnd1ClassName> <assoc_end1_name> in <AssocEnd2ClassName> <assoc_end2_name>
<i>query:</i>	yes
<i>exceptions:</i>	MofError (Invalid Object, Nil Object, Inaccessible Object)

The parameters to the “exists” operation are a pair of Instance values of the appropriate type for the Association. Since MOF link relationships are implicitly directional, the order of the parameters is significant.

“Invalid Object”, “Nil Object” and “Inaccessible Object” occurs if either of the parameters is a non-existent, nil or inaccessible Instance object.

<association_end1_name>

The “<association_end1_name>” operation queries the Instance object or objects that are related to a particular Instance object by a link in the current M1-level Association extent. When “isNavigable” is set to false for the AssociationEnd, the “<association_end1_name>” operation is suppressed.

<i>reflective analog:</i>	ref_query(<assoc_end1_designator>, <assoc_end1_name>)
<i>return type:</i>	<AssocEnd2ClassName>[<CollectionType>]
<i>parameters:</i>	in <AssocEnd1ClassName> <assoc_end1_name>
<i>query:</i>	yes
<i>exceptions:</i>	MofError (Invalid Object, Nil Object, Inaccessible Object, Underflow)

The <association_end1_name> parameter is the Instance object from which the caller wants to “navigate.” “Invalid Object,” “Nil Object,” and “Inaccessible Object” occur when the parameter is a non-existent, nil object or inaccessible Instance object.

The result type of the operation depends on the multiplicity of <AssociationEnd2>. If it has bounds of [1..1], the result type is the Instance type corresponding to the AssociationEnd’s “type.” Otherwise, it is a collection of the same Instance type, as described in Section 5.7.1.6, “Literal String Values,” on page 5-42.

“Underflow” occurs when <AssociationEnd2> has bounds [1..1] and the Instance object given by the parameter is not related in the current Association extent. It should not occur in other cases where the result type is a collection type. (If there is a multiplicity underflow, it is signalled by returning a collection value with too few elements as opposed to raising an exception.)

<association_end2_name>

This operation is the equivalent of “<association_end1_name>,” with the “end1” and “end2” interchanged.

add

The “add” operation creates a link in this Association between a pair of Instance objects. When “isChangeable” is set to false for either of the M2-level Association’s AssociationEnd, the “add” operation is suppressed.

<i>reflective analog:</i>	ref_add_link(Link{<assoc_end1_name>, <assoc_end2_name>})
<i>return type:</i>	none
<i>parameters:</i>	in <AssocEnd1ClassName> <assoc_end1_name> in <AssocEnd2ClassName> <assoc_end2_name>

<i>exceptions:</i>	MofError (Invalid Object, Nil Object, Inaccessible Object, Overflow, Duplicate, Reference Closure, Composition Closure, Composition Cycle)
--------------------	--

The two parameters to the “add” operation give the Instance objects at the two ends of the new link. “Invalid Object,” “Nil Object,” and “Inaccessible Object” occur if either of the parameter values is a non-existent, nil or inaccessible Instance object.

If one or other end of the Association has “isOrdered” set to true, the new link must be added so that it is the last member of the projection for the ordered AssociationEnd. The operation must also preserve ordering of the existing members of the ordered projection.

“Overflow” occurs when adding the new link would cause the size of the projection of either the first or second parameter object to exceed the upper bound for the opposite AssociationEnd. “Duplicate” occurs when the link set for the current Association extent already contains the link whose creation is requested.

“Reference Closure” occurs when either (or both) of the AssociationEnds has References, and the corresponding Instance object parameter does not belong to the same outermost Package extent as the Association object; see “The Reference Closure Rule” on page 4-17.

“Composition Closure” occurs when either AssociationEnd has “aggregation” set to “composite”, and either of the Instance object parameters does not belong to the same outermost Package extent as this Association object; see Section 4.9.2, “The Composition Closure Rule,” on page 4-20.

“Composition Cycle” occurs when adding the new link would create a cycle of composite / component relationships such that one of the Instance object parameters is a (ultimately) component of itself; see Section 4.8.2, “Aggregation “composite”,” on page 4-18.

add_before_<association_end1_name>

The “add_before_<association_end1_name>” operation creates a link between a pair of Instance objects at a given place in this Association. This operation is only generated when “isChangeable” a true for both AssociationEnds, and when the first AssociationEnd is multi-valued, ordered and navigable.

<i>reflective analog:</i>	ref_add_link_before(Link{<assoc_end1_name>, <assoc_end2_name>}, <assoc_end1_designator>, before); (See Section , “Reflective::RefAssociation,” on page 6-23).
<i>return type:</i>	none

<i>parameters:</i>	in <AssocEnd1ClassName> <assoc_end1_name> in <AssocEnd2ClassName> <assoc_end2_name> in <AssociationEnd1ClassName> before
<i>exceptions:</i>	NotFound, MofError (Invalid Object, Nil Object, Inaccessible Object, Overflow, Duplicate, Reference Closure, Composition Closure, Composition Cycle)

The first two parameters to the “add_before_<association_end1_name>” operation give the Instance objects at the two ends of the new link. “Invalid Object”, “Nil Object” and “Inaccessible Object” occur if either of the parameter values is a non-existent, nil or inaccessible Instance object.

The third parameter (“before”) gives an Instance object that determines the point at which the new link is inserted. “Invalid Object”, “Nil Object” and “Inaccessible Object” also apply to the “before” parameter value.

The “before” value should be present in the projection of the “<assoc_end2_name>” parameter value. If this so, the insertion point for the new link is immediately before the “before” value, otherwise the “NotFound” error occurs.

“Overflow,” “Duplicate,” “Reference Closure,” “Composition Closure,” and “Composition Cycle” occur as described for the “add” operation above.

add_before_<association_end2_name>

This operation is the equivalent of “add_before_<association_end1_name>,” with the “end1” and “end2” interchanged.

Note – The preconditions for generating the “add_before_<association_end1_name>” and “add_before_<association_end2_name>” operations are such that at most one of them may appear in an Association interface.

modify_<association_end1_name>

The “modify_<association_end1_name>” operation updates a link between a pair of Instance objects, replacing the Instance at AssociationEnd1 with a new Instance object. When AssociationEnd1 has “isChangeable” or “isNavigable” set to false, this operation is suppressed.

<i>reflective analog:</i>	ref_modify_link(Link{<assoc_end1_name>, <assoc_end2_name>}, <assoc_end1_designator> new_<assoc_end1_name>); (see Section , “Reflective::RefAssociation,” on page 6-23).
<i>return type:</i>	none

<i>parameters:</i>	in <AssocEnd1ClassName> <assoc_end1_name> in <AssocEnd2ClassName> <assoc_end2_name> in <AssocEnd2ClassName> new_<assoc_end1_name>
<i>exceptions:</i>	NotFound, MofError (Invalid Object, Nil Object, Inaccessible Object, Overflow, Underflow, Duplicate, Reference Closure, Composition Closure, Composition Cycle)

The first two parameters to the “modify_<association_end1_name>” operation should give the Instance objects at the ends of an existing link. “Invalid Object,” “Nil Object,” and “Inaccessible Object” occur if either of the parameter values are non-existent, nil or inaccessible Instance objects. “NotFound” occurs if the link does not exist in the current extent.

The third parameter (“new_<assoc_end1_name>”) gives the Instance object that is to replace the Instance at AssociationEnd1 for the selected link. “Invalid Object,” “Nil Object,” and “Inaccessible Object” also occurs if this parameter’s value is a non-existent, nil or inaccessible Instance object.

If the “<assoc_end1_name>” and “new_<assoc_end1_name>” parameters give the same Instance object, this operation is required to have no effect on the Association’s link set.

Note – The following error conditions apply to the state of the M1-level Association after the completion of the operation, not to any intermediate states.

“Underflow” occurs if completion of the operation would leave the M1-level Association in a state where

$$\text{size}(\text{Projection}(\text{<assoc_end1_name>})) \text{ less than } \text{<AssocEnd2>}.lower$$

“Overflow” occurs if completion of the operation would leave the M1-level Association in a state where

$$\text{size}(\text{Projection}(\text{new_<assoc_end1_name>})) \text{ greater than } \text{<AssocEnd2>}.upper$$

Note that the “Underflow” condition for the “new_<assoc_end1_name>” Instance should be treated as a deferred constraint.

“Duplicate” occurs if the operation would create a duplicate link in this M1-level Association extent. Similarly, “Composition Cycle” occurs if the operation creates a link that (on completion of the operation) would make the “<assoc_end2_name>” or “new_<assoc_end1_name>” objects components of themselves.

“Reference Closure” and “Composition Closure” occur if the operation would create a link that violates the corresponding closure rules; see Section 4.9.1, “The Reference Closure Rule,” on page 4-19 and Section 4.9.2, “The Composition Closure Rule,” on page 4-20.

If either AssociationEnd has “isOrdered” set to true, this operation must preserve ordering of the remaining members in the relevant projections of the ordered end. In addition:

- If AssociationEnd1 is ordered, the projection of “<assoc_end2_name>” must have “new_<assoc_end1_name>” in the position taken by “<assoc_end1_name>”.
- If AssociationEnd2 is ordered, the projection of “new_<assoc_end1_name>” must have “<assoc_end2_name>” as the last member.

modify_<association_end2_name>

This operation is the equivalent of “modify_<association_end1_name>”, with the “end1” and “end2” interchanged.

remove

The “remove” operation removes a link between a pair of Instance objects in the current Association extent. When either AssociationEnd or AssociationEnd2 has “isChangeable” set to false, the “remove” operation is suppressed.

<i>reflective analog:</i>	ref_remove_link(Link{<assoc_end1_name>, <assoc_end2_name>}); (see Section , “Reflective::RefAssociation,” on page 6-23).
<i>return type:</i>	none
<i>parameters:</i>	in <AssocEnd1ClassName> <assoc_end1_name> in <AssocEnd2ClassName> <assoc_end2_name>
<i>exceptions:</i>	NotFound, MofError (Nil Object, Underflow)

The two parameters to this operation give the Instance objects at both ends of the link that is to be removed from the current Association object’s link set. “Nil Object” occurs if either parameter value is a nil object reference.

“NotFound” occurs if the link to be deleted does not exist in the current Association extent.

Note – “Invalid Object” and “Inaccessible Object” does occur here. The “remove” operation needs to be capable of deleting links that involve Instance objects that have been deleted or that are inaccessible In the latter case, an implementation can usually fall back on local comparison of object references. If that fails (e.g., because there are multiple references for an Instance object) the implementation will typically be unable to distinguish the case from “NotFound”.

“Underflow” occurs if deleting the link would cause the size of the projection of either the “<assoc_end1_name>” or “<assoc_end2_name>” parameter value to be less than the corresponding “lower” bound.

If either AssociationEnd1 or AssociationEnd2 has “isOrdered” set to true, the “remove” operation must preserve the ordering of the remaining members of the corresponding projection.

5.8.11 Attribute Template

The Attribute Template defines the generation rules for M2-level Attributes whose “visibility” is “public_vis.” The Attribute Template declares operations to query and update the value of an Attribute. These operations appear on different interfaces, depending on the Attribute’s “scope”:

- IDL operations for instance-scoped Attributes appear in the Instance (“<ClassName>”) interface for the Attribute’s containing Class.
- IDL operations for classifier-scoped Attributes appear in the Class Proxy (“<ClassName>Class”) interface for the Attribute’s containing Class, and are inherited by the Instance interface.

The operations generated for an Attribute and their signatures depend heavily on the Attribute’s properties. For the purposes of defining the generated IDL, Attribute multiplicities fall into three groups:

- single-valued Attributes: multiplicity bounds are [1..1],
- optional-valued Attributes: multiplicity bounds are [0..1], and
- multi-valued Attributes: any other multiplicity.

Template

```
// if Attribute visibility is private or protected no IDL
// is generated
<<ANNOTATION TEMPLATE>>

// if lower = 0 and upper = 1
<AttributeType> <attribute_name> ()
    raises (Reflective::NotSet, Reflective::MofError);

// if lower = 1 and upper = 1
<AttributeType> <attribute_name> ()
    raises (Reflective::MofError);

// if upper > 1
<AttributeType><CollectionKind> <attribute_name> ()
    raises (Reflective::MofError);

// if upper = 1 and is_changeable
void set_<attribute_name> (in <AttributeType> new_value)
```



```

        raises (Reflective::MofError);

// if upper > 1 and is_changeable
void set_<attribute_name> (
    in <AttributeType><CollectionKind> new_value)
    raises (Reflective::MofError);
// if lower = 0 and upper = 1 and is_changeable
void unset_<attribute_name> ()
    raises (Reflective::MofError);
// if upper > 1 and is_changeable
void add_<attribute_name> (in <AttributeType> new_element)
    raises (Reflective::MofError);
// if upper > 1 and is_changeable and is_ordered
void add_<attribute_name>_before (
    in <AttributeType> new_element,
    in <AttributeType> before_element)
    raises (Reflective::NotFound, Reflective::MofError);

// if upper > 1 and is_changeable and is_ordered and not is_unique
void add_<attribute_name>_at (
    in <AttributeType> new_element,
    in unsigned long position)
    raises (Reflective::BadPosition, Reflective::MofError);
// if upper > 1 and is_changeable
void modify_<attribute_name> (
    in <AttributeType> old_element,
    in <AttributeType> new_element)
    raises (Reflective::NotFound, Reflective::MofError);
// if upper > 1 and is_changeable and is_ordered and not is_unique
void modify_<attribute_name>_at (
    in <AttributeType> new_element,
    in unsigned long position)
    raises (Reflective::BadPosition, Reflective::MofError);
// if upper > 1 and upper != lower and is_changeable
void remove_<attribute_name> (
    in <AttributeType> old_element)
    raises (Reflective::NotFound, Reflective::MofError);
// if upper > 1 and upper != lower and is_changeable and
// is_ordered and not is_unique
void remove_<attribute_name>_at (in unsigned long position)
    raises (Reflective::BadPosition, Reflective::MofError);

```

Operations

<attribute_name>

The “<attribute_name>” operation returns the value of the named Attribute.

<i>reflective analog:</i>	ref_value(<attribute_designator>); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	[0..1] - <AttributeType> [1..1] - <AttributeType> other - <AttributeType><CollectionKind>
<i>parameters:</i>	none
<i>query:</i>	yes
<i>exceptions:</i>	[0..1] - Unset, MofError [1..1] - MofError other - MofError

The signature of the “<attribute_name>” operation depends on the Attribute’s multiplicity as indicated above. Its behavior is as follows:

- In the [0..1] case, the operation either returns the Attribute’s optional value, or raises the NotSet exception to indicate that the optional value is not present.
- In the [1..1] case, the operation simply returns the Attribute’s single value.
- In other cases, the operations returns the Attribute’s collection value. In the case where the collection is empty the result value will be a sequence with length zero. No exception is raised in this case.

If the Attribute is instance-scoped, the operation is only available on Instance objects, and invoking it returns a value that is related to this Instance object. If the Attribute is classifier-scoped the operation can be invoked on both Class Proxy and Instance objects. In both cases, the operation returns a value that is related to all Instances for the Attribute’s Class in the current extent. For a more detailed comparison of classifier versus instance-scoped Attributes, see Section 4.4, “Semantics of Attributes,” on page 4-4.

The **MofError** exception may be raised to signal meta-model defined Constraint errors and implementation specific Semantic errors. However, an implementation generally should avoid doing this, for the reasons given in Section 4.11.6, “Access operations should avoid raising Constraint exceptions,” on page 4-24.

set_<attribute_name>

The “set_<attribute_name>” operation sets the value of the named Attribute.

<i>reflective analog:</i>	ref_set_value(<attribute_designator>, new_value); (see Section , “Reflective::RefObject,” on page 6-9).
---------------------------	--

<i>return type:</i>	none
<i>parameters:</i>	[0..1] - in <AttributeType> new_value [1..1] - in <AttributeType> new_value other - in <AttributeType><CollectionKind> new_value
<i>exceptions:</i>	[0..1] - MofError (Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle) [1..1] - MofError (Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle) other - MofError (Overflow, Underflow, Duplicate, Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle)

The signature of the “set_<attribute_name>” operation depends on the Attribute’s multiplicity as indicated above. Its behavior is as follows:

- In the single and optional-valued cases, the operation assigns the single value given by “new_value” to the named Attribute.
- In the multi-valued case, the operation assigns the collection value given by “new_value” parameter to the named Attribute.

When the Attribute has a lower bound of 0, its value can legally be empty:

- In the optional-valued case, the Attribute’s value is set to “empty” by invoking the “unset_<attribute_name>” operation described below.
- In the [0..N] case (where N is not 1), the Attribute’s value is set to empty by invoking “set_<attribute_name>” with a zero length sequence as the parameter.

“Composition Closure” and “Composition Cycle” are only possible when the type of the Attribute is a Class, and the Attribute has “composite” aggregation semantics:

- “Composition Closure” occurs when “new_value” or one of its members (in the multi-valued case) belongs to a different outermost Package extent to this object.
- “Composition Cycle” occurs when the operation would result in this object having itself as a direct or indirect component.

“Overflow,” “Underflow,” and “Duplicate” can only occur in the case of a multi-valued Attribute:

- “Overflow” occurs if the number of members in the “new_value” collection is greater than the Attribute’s upper bound.
- “Underflow” occurs if the number of members in the “new_value” collection is less than the Attribute’s lower bound.
- “Duplicate” occurs if the Attribute has “isUnique” set to true and the “new_value” collection contains duplicate values.

“Invalid Object” and “Inaccessible Object” occur when some Instance object is found to be non-existent or inaccessible. An implementation should only signal one these conditions when it prevents other consistency checking (e.g., testing for composition cycles).

unset_<attribute_name>

The “unset_<attribute_name>” operation sets the value of an optional-valued Attribute to empty. This operation is suppressed in the single-valued and multi-valued cases.

<i>reflective analog:</i>	ref_unset_value(<attribute_designator>); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	none
<i>exceptions:</i>	MofError

The “unset_<attribute_name>” operation is the only way to update an optional-valued Attribute to the “empty” state.

The MofError exception may be raised to signal meta-model defined Constraint errors and implementation specific Semantic errors.

add_<attribute_name>

The “add_<attribute_name>” operation updates a multi-valued Attribute by adding a new member value to its collection value. This operation is suppressed for optional and single-valued Attributes and for Attributes with “isChangeable” set to false.

<i>reflective analog:</i>	ref_add_value(<attribute_designator>, new_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <AttributeType> new_element
<i>exceptions:</i>	MofError (Overflow, Duplicate, Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle)

The “add_<attribute_name>” operation adds “new_element” to the collection of a changeable multi-valued Attribute. If the Attribute’s “multiplicity” has “isOrdered” set to true, the “new_element” is added after that current last element of the collection.

“Overflow” occurs if adding another element to the collection makes the number of elements it contains greater than the Attribute’s upper bound.

“Duplicate” occurs if the Attribute’s “multiplicity” has “isOrdered” set to true, and the “new_element” value is equal to an element of the Attribute’s current value.

“Composition Closure” and “Composition Cycle” are only possible when the type of the Attribute is a Class, and the Attribute has “composite” aggregation semantics:

- “Composition Closure” occurs when “new_element” belongs to a different outermost Package extent to this object.
- “Composition Cycle” occurs when the operation would result in this object being a direct or indirect component of itself.

“Invalid Object” and “Inaccessible Object” occur when some Instance Object is found to be non-existent or inaccessible. An implementation should only signal one these conditions when it prevents other consistency checking (e.g., testing for composition cycles).

add_<attribute_name>_before

The “add_<attribute_name>_before” operation updates a multi-valued Attribute by adding a new element at given position in its current collection value. The operation is suppressed for optional and single-valued Attributes, and for Attributes with “isChangeable” or “isOrdered” set to false.

<i>reflective analog:</i>	ref_add_value_before(<attribute_designator>, new_element, before_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <AttributeType> new_element in <AttributeType> before_element
<i>exceptions:</i>	NotFound, MofError (Overflow, Duplicate, Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle)

The “add_<attribute_name>_before” operation adds “new_element” to the collection at a given place within the collection value of an ordered, changeable, multi-valued Attribute. Insertion is required to preserve the initial order of the collection’s elements.

The “new_element” is inserted before the first occurrence in the Attribute’s collection of the value supplied as the “before_element” parameter (i.e., the occurrence with the smallest index). “NotFound” occurs when the “before_element” value is not present in the collection.

“Overflow,” “Duplicate,” “Composition Closure,” and “Composition Cycle” occur in equivalent situations to those for “add_<attribute_name>” above.

“Invalid Object” and “Inaccessible Object” occur when some Instance object is found to be non-existent or inaccessible. An implementation should only signal one these conditions when it prevents other consistency checking (e.g., testing for composition cycles).

add_<attribute_name>_at

The “add_<attribute_name>_at” operation updates a multi-valued Attribute by adding a new element at given position in its current collection value. It is provided for non-unique Attributes where an insertion point must be specified using an index. The operation is suppressed for optional and single-valued Attributes, for Attributes with “isChangeable” or “isOrdered” set to false, and for Attributes with “isUnique” set to true.

<i>reflective analog:</i>	ref_add_value_at(<attribute_designator>, new_element, position); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <AttributeType> new_element in unsigned long position
<i>exceptions:</i>	BadPosition, MofError (Overflow, Duplicate, Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle)

The “add_<attribute_name>_at” operation adds “new_element” at a given point within the collection value of an ordered, non-unique changeable, multi-valued Attribute. Insertion is required to preserve the initial order of the collection’s elements.

The insertion point is given by the value of the “position” parameter. This is the index of the collection member before which “new_element” should be inserted, with zero being the index for the first member. A “position” value equal to the current number of elements means that “new_element” should be added to the end of the collection. “BadPosition” occurs when the “position” value is greater than the number of elements in the collection. (It is not possible to create a collection value with “gaps” by adding elements with “position” values larger than the collection size.)

“Overflow,” “Duplicate,” “Composition Closure,” and “Composition Cycle” occur in equivalent situations to those described for “add_<attribute_name>” above.

“Invalid Object” and “Inaccessible Object” occur when some Instance object is found to be non-existent or inaccessible. An implementation should only signal one these conditions when it prevents other consistency checking (e.g., testing for composition cycles).

modify_<attribute_name>

The “modify_<attribute_name>” operation updates a multi-valued Attribute by replacing an existing member of its collection value. This operation is suppressed for optional and single-valued Attributes and for Attributes with “isChangeable” set to false.

<i>reflective analog:</i>	ref_modify_value(<attribute_designator>, old_element, new_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <AttributeType> old_element in <AttributeType> new_element
<i>exceptions:</i>	NotFound, MofError (Duplicate, Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle)

The “modify_<attribute_name>” operation replaces an occurrence of the value passed in the “old_element” parameter with the value of “new_element.” “NotFound” occurs if the “old_element” value is not present in the Attribute’s initial collection value.

If the Attribute has “isOrdered” set to true, the operation is required to preserve the initial order of the collection’s elements. If it also has “isUnique” set to false, then the operation is defined to replace the first occurrence (i.e., the one with the smallest index).

“Duplicate,” “Composition Closure,” and “Composition Cycle” occur in similar situations to those described for “add_<attribute_name>” above.

“Invalid Object” and “Inaccessible Object” occur when some Instance object is found to be non-existent or inaccessible. An implementation should only signal one these conditions when it prevents other consistency checking (e.g., testing for composition cycles).

modify_<attribute_name>_at

The “modify_<attribute_name>_at” operation updates a multi-valued Attribute by replacing a member of its collection value at a given position. It is provided for non-unique Attributes where the member to be modified must be specified using an index. This operation is suppressed for optional and single-valued Attributes and for Attributes with “isChangeable” set to false.

<i>reflective analog:</i>	ref_modify_value_at(<attribute_designator>, new_element, position); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none

<i>parameters:</i>	in <i><AttributeType></i> new_element in unsigned long position
<i>exceptions:</i>	BadPosition, MofError (Duplicate, Invalid Object, Inaccessible Object, Composition Closure, Composition Cycle)

The “modify_<attribute_name>_at” operation replaces the value whose index in the collection is given by the “position” parameter. “BadPosition” occurs if the “position” parameter is greater than or equal to the number of elements in the Attribute collection.

The replacement value is given by the “new_value” parameter. The operation is required to preserve the order of the collection’s elements.

“Duplicate,” “Composition Closure,” and “Composition Cycle” occur in similar situations to those described for “add_<attribute_name>” above.

“Invalid Object” and “Inaccessible Object” occur when some Instance object is found to be non-existent or inaccessible. An implementation should only signal one these conditions when it prevents other consistency checking (e.g., testing for composition cycles).

remove_<attribute_name>

The “remove_<attribute_name>” operation removes an existing member from a multi-valued Attribute. This operation is suppressed for optional and single-valued Attributes and for Attributes with “isChangeable” set to false. It is also suppressed when the lower and upper bounds are equal.

<i>reflective analog:</i>	ref_remove_value(<attribute_designator>, old_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <i><AttributeType></i> old_element
<i>exceptions:</i>	NotFound, MofError(Underflow)

The “remove_<attribute_name>” operation removes an occurrence of the value passed in the “old_element” parameter. “NotFound” occurs if the “old_element” value is not present in the Attribute’s collection value.

If the Attribute has “isOrdered” set to true, the operation is required to preserve the initial order of the collection’s elements. If it also has “isUnique” set to false, then the operation is defined to remove the first occurrence (i.e., the one with the smallest index).

“Underflow” occurs if removing an element makes the number of elements in the collection less than the Attribute’s lower bound.

Note – The “remove_<attribute_name>” operation should not signal an exception if it finds that some Instance object is non-existent or inaccessible. If the object in question is the object to be removed from the Attribute, it should be removed. Otherwise, the condition should be silently ignored.

remove_<attribute_name>_at

The “remove_<attribute_name>_at” operation removes the member at a given position from a multi-valued Attribute. This operation is suppressed for optional and single-valued Attributes, and for Attributes with “isChangeable” or “isOrdered” set to false or “isUnique” set to true. It is also suppressed when the lower and upper bounds are equal.

<i>reflective analog:</i>	ref_remove_value_at(<attribute_designator>, position); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in unsigned long position
<i>exceptions:</i>	BadPosition, MofError(Underflow)

The “remove_<attribute_name>_at” operation removes the element of an Attribute’s collection value whose (zero based) index is given by the “position” parameter. “BadPosition” occurs if the “position” value is greater than or equal to the number of elements in the Attribute’s collection value.

“Underflow” occurs if removing an element makes the number of elements in the collection less than the Attribute’s lower bound.

5.8.12 Reference Template

The Reference Template defines the IDL generation rules for a Reference whose “visibility” is “public_vis”. The IDL generated for a Reference is declared within the scope of <ClassName>Class interface definition. The IDL generated by the Reference Template provides the operations to return the value of the Reference as well as operations to modify it. The IDL generated is dependent upon the multiplicity, mutability, and ordering of the specified Reference.

The Reference Template defines the IDL generation rules for References. It declares operations on the Instance interface to query and update links in the Association object for the current extent.

The operations generated for a Reference and their signatures depend heavily on the properties of the referenced AssociationEnd which are also mirrored on the Reference itself. For the purposes of defining the generated IDL, Reference multiplicities fall into three groups:

- single-valued References: multiplicity bounds are [1..1],

- optional-valued References: multiplicity bounds are [0..1], and
- multi-valued References: any other multiplicity.

The generated operations for a Reference are designed to have similar signatures and behaviors to those for an instance-scoped Attribute with the same multiplicity and changeability settings.

Note – Recall that Reference is only well formed if the referenced AssociationEnd has “isNavigable” set to true. Similarly, a Reference’s “isChangeable” can only be true if the referenced AssociationEnd’s “isChangeable” is also true.

Template

*// If the Reference has visibility of protected or private, no IDL
// is generated*

<<ANNOTATION TEMPLATE>>

// operations to return the Reference value

// if lower = 0 and upper = 1

<ReferenceClass> <reference_name> ()

raises (Reflective::NotSet, Reflective::MofError);

// if lower = 1 and upper = 1

<ReferenceClass> <reference_name> ()

raises (Reflective::MofError);

// if upper > 1

<ReferenceClass><Multiplicity> <reference_name> ()

raises (Reflective::MofError);

// operations to modify the Reference value

// if upper = 1 and is_changeable

void set_<reference_name> (

in <ReferenceClass> new_value)

raises (Reflective::MofError);

// if upper > 1 and is_changeable

void set_<reference_name> (

in <ReferenceClass><Multiplicity> new_value)

raises (Reflective::MofError);

```

// if lower = 0 and upper = 1 and is_changeable
void unset_<reference_name> ()
    raises (Reflective::MofError);

// if upper > 1 and is_changeable
void add_<reference_name> (
    in <ReferenceClass> new_element)
    raises (Reflective::MofError);

// if upper > 1 and is_changeable and is_ordered
void add_<reference_name>_before (
    in <ReferenceClass> new_element,
    in <ReferenceClass> before_element)
    raises (Reflective::NotFound, Reflective::MofError);

// if upper > 1 and is_changeable
void modify_<reference_name> (
    in <ReferenceClass> old_element,
    in <ReferenceClass> new_element)
    raises (Reflective::NotFound, Reflective::MofError);

// if upper > 1 and lower != upper and is_changeable
void remove_<reference_name> (
    in <ReferenceClass> old_element)
    raises (Reflective::NotFound, Reflective::MofError);

```

<reference_name>

The “<reference_name>” operation reads the value of Reference. The signature of the operation depends on the multiplicity of the Reference.

<i>reflective analog:</i>	ref_value(<reference_designator>); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	[0..1] - <ReferenceClass> [1..1] - <ReferenceClass> other - <ReferenceClass><CollectionKind>
<i>parameters:</i>	none
<i>exceptions:</i>	[0..1] - NotSet, MofError [1..1] - MofError (Underflow) other - MofError

The “<reference_name>” operation’s signature is determined by the multiplicity of the Reference, and hence the referenced AssociationEnd, as shown above.

In each case, the operation calculates and returns the projection of “this” object in the link set of the referenced AssociationEnd’s Association for the current outermost extent:

- In the [0..1] case, the operation returns the projected Instance object if there is one, and raises the **Reflective::NotSet** exception if there is not.
- In the [1..1] case, the operation normally returns a single Instance object. However, if the projection contains no elements, this is signalled as a **Reflective::MofError** exception with “error_kind” of “Underflow”.
- In other cases, the operation returns the projection using a sequence value. If the projection is empty the result is a sequence of length zero. If it contains fewer elements than the Reference’s lower bound, those that it does contain are returned.

Note – Under no circumstances should the “<reference_name>” operation return a nil object reference or a sequence that includes a nil object reference.

set_<reference_name>

The “set_<reference_name>” operation assigns a new value to a Reference. The signature of the operation depends on the multiplicity of the Reference. If “isChangeable” is set to false for the Reference, this operation is suppressed.

<i>reflective analog:</i>	ref_set_value(<reference_designator>, new_value); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	[0..1] - in <ReferenceClass> new_value [1..1] - in <ReferenceClass> new_value other - in <ReferenceClass><CollectionKind> new_value
<i>exceptions:</i>	[0..1] - MofError (Overflow, Underflow, Invalid Object, Nil Object, Inaccessible Object, Composition Closure, Composition Cycle, Reference Closure) [1..1] - MofError (Overflow, Underflow, Invalid Object, Nil Object, Inaccessible Object, Composition Closure, Composition Cycle, Reference Closure) other - MofError (Overflow, Underflow, Duplicate, Nil Object, Inaccessible Object, Invalid Object, Composition Closure, Composition Cycle, Reference Closure)

The “set_<reference_name>” operation’s signature is determined by the multiplicity of the Reference, and hence the referenced AssociationEnd, as shown above.

In each case, the operation replaces the set of links in the extent for the referenced AssociationEnd’s Association. The behavior is as follows:

- In the [0..1] and [1..1] case, the caller passes a single Instance object in the “new_value” parameter that is used to create the replacement link.
- In other cases, the “new_value” parameter is a sequence of Instance objects that are used to create the replacement links. If the sequence is empty, no replacement links will be created.

The projection for an optional-valued Reference can only be set to “empty” using the “unset_<reference_name>” operation; see below.

The ordering semantics of the “set_<reference_name>” operation depend on the setting of “isOrdered” in the “multiplicity” for the Reference’s “referencedEnd” and “exposedEnd” AssociationEnds:

- If neither of the AssociationEnds has “isOrdered” set to true, the Association has no ordering semantics.
- If the “referencedEnd” AssociationEnd has “isOrdered” set to true, the order of the elements of the projection of “this” Instance after the operation has completed must be the same as the order of the elements of the “new_value” parameter.
- If the “exposedEnd” AssociationEnd has “isOrdered” set to true, the order of the elements of the “new_value” parameter (if it is multi-valued) are irrelevant. Instead, the operation is required to preserve the ordering of the projections that contained “this” object, both before and after the update, as follows:
 - If “this” object is in a projection of some other Instance object before the operation but not afterwards, the order of the projection must be preserved, with “this” object removed.
 - If “this” object is in a projection of some other Instance object after the operation but not before, the order of the projection must be preserved, *and* “this” object must be added at the end of the projection.
 - If “this” object is in a projection of some other Instance object both before *and* after the operation, the “before” and “after” versions of the projection must be identical.
- It is impossible for both of the AssociationEnds to have “isOrdered” set to true.

A large number of error conditions can occur, depending on “new_value”, the current state of the Association and the multiplicity of the Reference’s “referencedEnd” and “exposedEnd” AssociationEnds:

- “Invalid Object,” “Nil Object,” and “Inaccessible Object” occur if any of the supplied Instance objects is a non-existent, nil or inaccessible Instance object.

- “Overflow” can occur in two cases. First, it occurs when the “new_value” parameter contains more elements than is allowed by the “referencedEnd”’s upper bound. Second, it occurs when the projection of an element of “new_value” after completion of the operation would have more elements than is allowed by the “exposedEnd”’s upper bound.
- “Duplicate” occurs for a multi-valued Reference when the “new_value” parameter collection contains two or more occurrences of the same Instance object.
- “Underflow” can also occur in two cases. First it occurs when the “new_value” parameter contains fewer elements that is allowed by the “referencedEnd”’s lower bound. Second, it occurs when the projection of an element of “new_value” after completion of the operation would have fewer elements than is allowed by the “exposedEnd”’s lower bound, *and* fewer elements than it had before the operation commenced.
- “Reference Closure” occurs when “new_value” (in the [0..1], [1..1] case) or one of its elements (in the “other” case) belongs in a different outermost extent to “this” object.
- “Composition Closure” occurs in the same situation as “Reference Closure,” where the referenced Association has composite aggregation semantics.
- “Composition Cycle” occurs when the referenced Association has composite aggregation semantics, and the update would make “this” object a component of itself.

unset_<reference_name>

The “unset_<reference_name>” operation sets an optional-valued Reference to empty. If “isChangeable” is set to false for the Reference, or if the bounds are not [0..1], this operation is suppressed.

<i>reflective analog:</i>	ref_reset_value(<reference_designator>); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	none
<i>exceptions:</i>	MofError

The “unset_<reference_name>” operation removes the link for this object from the link set of the referenced Association, should it exist. If no such link exists, the operation does nothing.

If the “exposedEnd” AssociationEnd has “isOrdered” set to true, the operation preserves the ordering of the projection that initially contains “this” Instance object.

add_<reference_name>

The “add_<reference_name>” operation adds an Instance object to a Reference collection. If “isChangeable” is set to false for the Reference, or the Reference’s upper bound is 1, this operation is suppressed.	
<i>reflective analog:</i>	ref_add_value(<reference_designator>, new_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <ReferenceClass> new_element
<i>exceptions:</i>	MofError (Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle)

The “add_<reference_name>” operation adds the “new_element” Instance to a multi-valued Reference collection by creating a link in the corresponding Association’s link set. “Invalid Object,” “Nil Object,” or “Inaccessible Object” occur if the “new_element” parameter is a non-existent, nil or inaccessible Instance object.

If the “referencedEnd” AssociationEnd has “isOrdered” set to true, the new link should be created so that “new_element” is the last element of the projection of “this” object. Alternatively, if the “exposedEnd” AssociationEnd has “isOrdered” set to true, the new link should be created so that “this” object is the last element of the projection of the “new_element” object. In either case, the operation should preserve the order of other elements in the respective ordered projections.

“Overflow” occurs if the number of elements in the projections of either the “this” object or the “new_element” object already equals the respective AssociationEnd’s upper bound.

“Duplicate” occurs if the operation would create a duplicate link in the link set for the referenced Association. For example, when the “new_element” value is a duplicate of a value in the current Reference collection.

“Reference Closure,” “Composition Closure,” and “Composition Cycle” all occur in similar situations to those described above for the “set_<reference_name>” operation.

add_<reference_name>_before

The “add_<reference_name>_before” operation adds an Instance object at a particular place in an ordered Reference collection. If “isChangeable” or “isOrdered” is set to false for the Reference, this operation is suppressed.	
<i>reflective analog:</i>	ref_add_value_before(<reference_designator>, new_element, before_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <ReferenceClass> new_element in <ReferenceClass> before_element
<i>exceptions:</i>	NotFound, MofError (Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle)

The “add_<reference_name>_before” operation is a more specialized version of the “add_<reference_name>” operation described previously. It creates a link between “this” object and the “new_element” Instance object so that it appears in a designated place in “this” object’s projection.

The “before_element” parameter gives the Instance object in the projection of “this” before which the “new_element” object should be inserted. “Invalid Object,” “Nil Object,” and “Inaccessible Object” occur if either “new_element” or “before_element” is a non-existent, nil or inaccessible Instance object. “Not Found” occurs if “before_element” is not present in the projection of “this” object.

The new link is created such that the “new_element” object appears immediately before the “before_element” value in the projection of “this” object. Apart from this, the order of the projection’s elements is unchanged.

“Overflow,” “Duplicate,” “Reference Closure,” “Composition Closure” and “Composition Cycle” all occur in equivalent situations to those described above for the “add_<reference_name>” and “set_<reference_name>” operations.

modify_<reference_name>

The “modify_<reference_name>” operation updates a Reference collection, replacing one element with another. If the Reference is not multi-valued or its “isChangeable” multiplicity flag is set to false, this operation is suppressed.

<i>reflective analog:</i>	ref_modify_value(<reference_designator>, old_element, new_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none

<i>parameters:</i>	in <ReferenceClass> old_element in <ReferenceClass> new_element
<i>exceptions:</i>	NotFound, MofError (Underflow, Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle)

The “modify_<reference_name>” operation updates the link set so that the projection of “this” object has “new_element” in place of “old_element.” The operation is notionally equivalent to either

```
<the_association>.modify_<association_end1>(
    old_element, <this>, new_element)
```

or

```
<the_association>.modify_<association_end2>(
    <this>, old_element, new_element)
```

where <the_association> is the current outermost extent’s M1-level Association object for the referenced M2-level Association.

The “old_element” and “new_element” parameters must both give usable Instance objects. “Invalid Object”, “Nil Object” or “Inaccessible Object” occur if either is a non-existent, nil or inaccessible object.

The “old_element” object must be an element of the projection of “this” object; i.e. a link must already exist between “this” and “old_element”. “NotFound” occurs if this is not the case. If “old_element” and “new_element” are the same Instance object, the operation is required to do nothing at all.

If the referenced Association is ordered, the operation is required to preserve ordering as follows:

- If the “referencedEnd” AssociationEnd has “isOrdered” set to true, the order of the elements in the projection of “this” object should be preserved, with “new_element” occupying the same position as “old_element” did before the update.
- If the “exposedEnd” AssociationEnd has “isOrdered” set to true, the order of the elements in the projections of “old_element” and “new_element” should be preserved, except that “this” is removed from the former projection and added to the end of the latter projection.

“Overflow” occurs when the number of elements in the projection of “new_element” would be greater than the upper bound for the “exposedEnd” AssociationEnd.

“Underflow” occur when the number of elements in the projection of “old_element” would be decreased, and it would be less than the lower bound of the “exposedEnd” AssociationEnd. (In the case where “old_element” and “new_element” are the same object, the operation does not alter the number of elements in the projection. Hence “Overflow” cannot be signalled, even if the number of elements is less than the bound.)

“Duplicate” occurs if the “modify_<reference_name>” operation would introduce a duplicate into the projection. (Care should be taken to avoid signalling “Duplicate” in the case where “old_element” and “new_element” are the same object.)

“Reference Closure”, “Composition Closure” and “Composition Cycle” all occur in equivalent situations to those described above for the “add_<reference_name>” and “set_<reference_name>” operations.

remove_<reference_name>

The “remove_<reference_name>” operation updates a Reference collection by removing an element. If the Reference is not multi-valued or its “isChangeable” multiplicity flag is set to false, this operation is suppressed. It is also suppressed if the Reference’s lower and upper bounds are equal.

<i>reflective analog:</i>	ref_remove_value(<reference_designator>, old_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	none
<i>parameters:</i>	in <ReferenceClass> old_element
<i>exceptions:</i>	NotFound, MofError (Underflow)

The “remove_<reference_name>” operation updates the link set (i.e. by removing a link) so that the projection of “this” object no longer contains “old_element”. “NotFound” occurs if there is no link to be deleted.

Note – The “remove_<reference_name>” operation should be able to cope with removal of a link when the object at the other end of a link is non-existent or inaccessible.

If the referenced Association is ordered, the operation is required to preserve the ordering of the projection with the ordered collection value.

“Underflow” occur when the number of elements in the projections of “old_element” and “this” would be less than the lower bounds of the respective AssociationEnds.

5.8.13 Operation Template

The Operation Template defines the IDL generation rules for M2-level Operations whose “visibility” is “public_vis.” It generates an IDL operation within the scope of an Instance or Class Proxy interface, depending on the scope of the M2-level Operation.

Template

```
// If the Operation has visibility of protected or private, no IDL
// is generated
```

<<ANNOTATION TEMPLATE>>

```
// if Operation contains no "return" Parameter
void <operation_name>{
// else
<ReturnParamType>[<CollectionKind>] <operation_name>{
    // for each contained "in", "out" or "inout" Parameter
    <direction> <ParamType>[<CollectionKind>] <param_name>, ...
    )
    raises (
        // for each Exception raised by the Operation
        <ExceptionName>, ... // (a trailing comma is required)
        Reflective::MofError);

// for each Constraint contained by this Operation
<<CONSTRAINT_TEMPLATE>>
```

<operation_name>

An “<operation_name>” operation invokes an implementation specific method to perform the behavior implied by the M2-level Operation model element.

<i>reflective analog:</i>	ref_invoke_operation(<reference_designator>, old_element); (see Section , “Reflective::RefObject,” on page 6-9).
<i>return type:</i>	no return param - void [0..1] return param - <ParamType>Bag <param_name> [1..1] return param - <ParamType> <param_name> other return param - <ParamType><CollectionKind> <param_name>
<i>parameters:</i>	<direction> <ParamType>[<CollectionKind>], ...
<i>exceptions:</i>	<ExceptionName>, ... MofError (Overflow, Underflow, Duplicate, Invalid Object)

An “<operation_name>” operation invokes an implementation specific method. While the behavior of the method itself is beyond the scope of the IDL mapping, the signature of the IDL operation is defined by the mapping, along with some parameter checking semantics.

The return type for an “<operation_name>” operation is generated from the M2-level Operation’s (optional) return Parameter. For example, the contained Parameter object whose “direction” attribute has the value “return_dir”. The return type is as follows:

- If there is no return Parameter, the return type is “void”.
- If the return Parameter has “multiplicity” bounds of “[1..1]”, the return type is the “type” of the Parameter; i.e. <ParameterType>.
- If the return Parameter some other “multiplicity” bounds, the return type is a collection type determined by the bounds. For example, <ParameterType><CollectionKind>, as described in Section 5.7.1.6, “Literal String Values,” on page 5-42.

The parameter declarations for an “<operation_name>” operation are generated from the M2-level Operation’s Parameter, excluding the return Parameter (if any). For each non-return Parameter of the Operation, in the defined order, the “<operation_name>” declaration has a parameter declaration consisting of the following:

- The “<direction>” is produced by rendering the Parameter’s “direction” as “in,” “out,” or “inout” as appropriate.
- The “<ParameterType>[<CollectionKind>]” is produced from the Parameter’s “type” and “multiplicity” as follows:
 - If the Parameter has “multiplicity” bounds of “[1..1]”, the <CollectionKind> is omitted.
 - If the Parameter has “multiplicity” bounds other than “[1..1]”, <CollectionKind> is generated according to Section 5.7.1.6, “Literal String Values,” on page 5-42.
- The “<parameter_name>” is produced by rendering the Parameter’s name.

The list of exceptions raised by an “<operation_name>” operation is generated from the M2-level Operation’s “exceptions”. The generated “raises” list consists of an appropriately qualified identifier for each M2-level Exception in the Operation’s “exceptions” list, followed by the qualified identifier for the MofError exception. The “raises” list should of course be comma separated as required by the syntax for OMG IDL.

While meta-model specific error conditions should be signalled by raising exceptions corresponding to the Operation’s “exceptions” list, MofError is used to signal the following structural errors relating to the values supplied by the caller for “in” and “inout” parameters.

- “Overflow” occurs when the supplied collection value for a multi-valued parameter has more elements than is allowed by the M2-level Parameter’s upper bound.
- “Underflow” occurs when the supplied collection value for a multi-valued parameter has fewer elements than is allowed by the M2-level Parameter’s lower bound.
- “Duplicate” occurs when a multi-valued M2-level Parameter has “isUnique” set to true, and the supplied collection value contains members that are equal according to the definitions in Section 4.9, “Closure Rules,” on page 4-18.

- “Invalid Object” can occur if an Instance object typed parameter value or element is a reference to a non-existent (i.e., deleted) or inaccessible object. This condition will occur if duplicate checking finds an Instance object that it cannot test for equality. It can also occur if the semantics of the Operation require an Instance object reference to be usable.

Like all other operations that have MofError in their signature, an “<operation_name>” operation can use MofError to signal Constraint errors and Semantic errors.

5.8.14 Exception Template

The Exception template defines the IDL generation rules for M2-level Exceptions whose “visibility” is “public_vis.”

Template

```
// If the Exception has visibility of protected or private, no IDL
// is generated
```

```
<<ANNOTATION TEMPLATE>>
exception <ExceptionName> {
    // for each Parameter
    <ParameterType>[<CollectionKind>] <parameter_name>; ...
};
```

Description

The generated IDL for an M2-level Exception is an IDL exception. The declaration appears within an IDL interface or module corresponding to the Exception’s M2-level container. In the case of an M2-level Class, this is the Class Proxy interface so that the IDL exception is available to be raised by classifier-scoped Operations.

The fields of the IDL exception are generated from the Exception’s Parameters in a way that is similar to Operation Parameters:

- An Exception Parameter whose multiplicity has a “[1..1]” bound is mapped to a field whose type is “<ParameterType>”.
- An Exception Parameter whose multiplicity has any other bound is mapped to a field whose type is of the form “<ParameterType><CollectionKind>,” generated according to the rules in Section 5.7.1.6, “Literal String Values,” on page 5-42.

5.8.15 Constant Template

The Constant Template defines the rules for generating IDL constant declarations from M2-level Constants.

Template

```
<<ANNOTATION TEMPLATE>>
```

```
const <ConstantType> <CONSTANT_NAME> = <CONSTANTVALUE>;
```

The generated IDL for an M2-level Constant is an IDL constant declaration. The IDL appears an interface or module corresponding to the Constant's M2-level container. In the container is a Class, the declaration appears within the Class Proxy interface.

The IDL generation process needs to produce a valid IDL literal value of the appropriate type from the Constant's "value."

5.8.16 DataType Template

The DataType Template defines the rules for generating IDL for an M2-level DataType whose "visibility" is "public_vis." This typically consists of an IDL type declaration for the data type, followed by one or more collection type declarations, as required.

Note – If the IDL mapping preconditions are strictly observed, the template will only generate IDL declarations for the DataType's type in cases where this is appropriate.

Template

```
// If the DataType's name does not a map to a valid IDL identifier,
// valid name, no IDL is generated
// If the DataType's visibility is protected or private, no IDL
// is generated
```

```
<<ANNOTATION TEMPLATE>>
```

```
// generate the DataType's type declaration
// if the DataType's typecode kind is tk_alias
typedef <TYPECODE.CONTENTS.TYPESPEC> <datatype_name>;
```

```
// else
//   the DataType's typecode kind is for a named IDL data type
//   (i.e. a struct, union or enumeration type)
<TYPECODE.TYPESPEC>;
```

```
// For each Constraint contained by this DataType
```

```
<<CONSTRAINT_TEMPLATE>>
```

```
// if collection types for the DataType are used within the
// current outermost Package, for each collection type:
typedef sequence < <DataTypeName> > <DataTypeName><CollectionKind>;
```

Description

A DataType template only generates IDL type declarations for named M2-level DataTypes that represent CORBA data types:

- A DataType whose name does not map to a valid identifier represents a use of an anonymous data type. This does not require an IDL data type declaration, or any collection type declarations at this point. If they are required, the collection type declarations will appear at the beginning of the outermost module.
- A DataType whose “typeCode” kind is “tk_objref” represents a use of either a Class or an externally defined CORBA interface. In either case, the DataType template generates nothing.

The generation process is effectively a “reverse compilation” of the DataType’s “typeCode” into OMG IDL text. The process is not spelled out in detail here, but can be inferred from the syntax of type declarations in IDL, the structure of TypeCodes and the restrictions that the MOF Model and the IDL mapping place on them.

If a DataType is used in an M2-level Attribute, Reference, AssociationEnd or Parameter with a multiplicity other than “[1..1]”, it is likely that the generated IDL for the current meta-model will contain a reference to one or more collection types for the DataType’s mapped type. If this is so, the template also generates sequence type declarations for the required collection types; see Section 5.7.2, “Generation Rules for Collection Types,” on page 5-42.

Note – In the interests of IDL footprint size, the DataType template should only generate collection type declarations if they are needed.

Finally, if a DataType contains any Constraints, the corresponding constraint name strings need to be generated.

5.8.17 Constraint Template

The Constraint template defines the rules for generating the requisite error kind string declaration for an M2-level Constraint.

Template

```
<<ANNOTATION TEMPLATE>>
const string <CONSTRAINT_NAME> = "<constraint.string>";
```

Description

The Constraint template generates an IDL string constant whose name is based on the M2-level Constraint name. If the Constraint is contained by an M2-level DataType or Operation, the constant declaration is generated within the scope of the Constraint

container's container. If this results in a name collision, the meta-modeler can solve the problem using a substitute name tag as described in Section 5.6.2.1, "Substitute Name," on page 5-37.

The "`<constraint.string>`" value is generated to match the following syntax (expressed in EBNF):

```
<constraint.string> ::= [ <IDL prefix> ] `:constraint.`
                        ( <container_name> `.` )+ <constraint_name>
```

The components of the error kind string value are as follows:

- If the meta-model has an IDL prefix (see Section 5.6.1.1, "IDL Prefix," on page 5-36), the string starts with the value of this prefix.
- Next there is a colon (":") to separate the prefix from the rest of the string.
- Next there is the fixed string "constraint" to indicate that the class of error, followed by a period ".").
- Next there are a series of Format 2 renderings of the names of the Constraint's enclosing containers. These are separated by period (".") characters, and followed another period.
- The value ends with the Format 2 rendering of the name of the Constraint itself.

5.8.18 Annotation Template

The Annotation template optionally generates IDL comments for an M2-level ModelElement's "annotation". This template should be regarded as indicative rather than normative.

Template

```
// Annotation comments may optionally be suppressed by the IDL
// generator

// Annotation comments may use the "/*...*/" style

/* <line 1 of the ANNOTATION>
   <line 2 of the ANNOTATION>
   . . .
   <line N of the ANNOTATION> */

// or the "//" style

// <line 1 of ANNOTATION>
// <line 2 of ANNOTATION>
// . . .
```

```
// <line N of the ANNOTATION>
```

Description

The Annotation template optionally includes the “annotation” for a ModelElement in the generated IDL as an IDL comment. It is anticipated that a vendor’s IDL generator would give some control over the way that these comments are generated. For example, allowing the user to

- suppressing the comments completely,
- choose between the two styles of comments, and
- choose whether or not to respect embedded line breaks and other markup.

Contents

This chapter contains the following topics.

Topic	Page
“Introduction”	6-1
“The Reflective Interfaces”	6-3
“The CORBA IDL for the Reflective Interfaces”	6-29

6.1 Introduction

One of the advantages of meta-objects (in the general sense) is that they allow a program to use objects without prior knowledge of the objects' interfaces. In the MOF context, an object's M2-level meta-object allows a program to “discover” the nature of any M1-level MOF object, both at a syntactic level and at a deeper level. With this information in hand, the MOF's Reflective interfaces allow a program to:

- create, update, access, navigate and invoke operations on M1-level Instance objects,
- query and update links using M1-level Association objects, and
- navigate an M1-level Package structure

without using meta-model specific interfaces.

Note – The functionality above is all available through the "model specific" interfaces defined by the IDL mapping described in this chapter. The Reflective interfaces do not allow a program to access or update MOF objects contrary to their meta-object descriptions. For example, they cannot be used to create, access or update Attributes that do not exist, or to bypass Constraint checking.

In addition, the Reflective interfaces allow the program to:

- find an M1-level object's M2-level meta-object,
- find a MOF object's container(s) and enclosing Package(s),
- test for MOF object identity, and
- delete a MOF object.

Note – While many of these capabilities are correctly described as reflective, the MOF does not offer the full repertoire of reflective programming features. Since it does not define object behavior, the MOF does not define interfaces for reflective behavior modification. Even if it did, these interfaces could not be implemented in many CORBA contexts.

The CORBA Interface Repository (IR) and the Dynamic Invocation Interface (DII), provide similar capabilities in the context of a CORBA object's Interface. However, using the IR and DII for this purpose means that the user cannot make use of the richer semantic information in models defined using the MOF meta-model. For example, the IR can tell the user that the "Model::Contains" IDL interface has an operation called "exists;" however, it is only by using MOF meta-objects that the user knows that the "exists" operation tests whether one object "contains" another one.

The MOF Reflective module contains four "abstract" interfaces that are inherited by the M1-level interfaces for a model that are generated from a meta-model by the IDL mapping.

1. The Reflective::RefObject interface provides common operations for M1-level Instance objects and Class Proxy objects.
2. The Reflective::RefAssociation interface provides common operations for M1-level Association objects.
3. The Reflective::RefPackage interface provides a common operations for M1-level Package objects.
4. The Reflective::RefBaseObject interface provides common operations for all MOF objects.

Since the M2-level interfaces for the MOF Model are generated by this means, they also inherit from the Reflective interfaces.

6.2 *The Reflective Interfaces*

This section describes the interfaces defined in the "Reflective" module. These interfaces are modeled on the interfaces that are produced by the IDL mapping. However, there are some important differences:

- Reflective operations pass the values of Attributes and References, and of the Parameters to Operations and Exceptions as CORBA Any values. The mapped versions of these operations pass the values using precise types according to the meta-model.
- Reflective operations on Associations pass Instance objects with the type RefObject. The mapped versions of these operations pass Instance objects using their true types.
- The "target" feature for a Reflective operation is passed as a "designator" parameter whose type is a MOF meta-object. In the mapped case, the target is implicit in the mapped operation name.

As stated previously, the Reflective versions of operations which are defined in the mapped IDL do not allow a program to violate the information and computational models implied by the meta-model definition. This includes not allowing operations that, while meaningful for a model, are not possible using the mapped interfaces. For example, while it might be meaningful to call "refSetValue" on an optional Attribute passing an "empty" argument (encoded appropriately), this is not allowed: the program must use "refSetValue".

This section consists of a subsection that explains some common patterns that are used for encoding parameters used by many Reflective operations. The remaining four subsections describe each Reflective interface in turn.

6.2.1 *Reflective Argument Encoding Patterns.*

The Reflective module make heavy use of the CORBA Any type to provide meta-model independent interfaces. This section defines some common patterns used throughout the Reflective interfaces for encoding parameter values in Anys.

Note – It is important that the type information (expressed as CORBA TypeCodes) in the encoded Anys be precisely as specified below. In particular, collection type aliases and their names are mandatory.

If the base type of the value-defining feature is a DataType, the TypeCode in the encoded Any must be the full TypeCode for the base type. Type aliases must not be optimized away, and all optional names (e.g., of struct types, fields, and so on) must be present. (Optimization of type information in Anys should done at the ORB level if at all.)

Note – The IDL templates can cause multiple copies of a collection type to be declared in the generated IDL for a composite meta-model. Since the copies are logically identical, MOF client and server code needs to take care when extracting collection values from Anys. In particular, if the stub-generated extraction operations will fail if the "wrong" copy of a collection type is used, the code may need to use DynAny instead.

6.2.1.1 The Standard Value Encoding Pattern

This pattern is used for encoding complete values as ValueTypes. It is used in most cases where a reflective operation requires or provides a complete value for an element that may be collection valued (depending on the multiplicity). Examples that use this pattern are values for Operation arguments and results, values for Exception fields and Attribute initial values in a create operation.

Table 6-1 Standard Value Encoding Pattern

Bounds	ValueType Encoding	Notes
[0..1]	Any(alias(seq(<type>, 0))) where the alias name is <typeName>Bag	An "optional" feature value with no elements is encoded as zero length sequence.
[1..1]	Any(<type>)	
others	Any(alias(seq(<type>, 0))) where the alias name is <typeName><CollectionKind>	A "multi-valued" feature value with no elements is encoded as a zero length sequence.

6.2.1.2 The Alternate Value Encoding Pattern

The standard pattern for encoding complete values (above) does not fit well with the IDL templates for the specific "get" and "set" operations. To improve the alignment between the reflective and specific interfaces, the following alternative pattern is used for the “The “The operation fetches the current value of the Attribute or Reference denoted by the "feature" argument. If this object is a Class proxy, only classifier scoped Attributes can be fetched.” operation fetches the current value of the Attribute or Reference denoted by the "feature" argument. If this object is a Class proxy, only classifier scoped Attributes can be fetched.” and “The “refSetValue” operation assigns a new value to an Attribute or Reference for an object. The assigned value must be a single value or a collection value depending on the feature’s multiplicity.” operations for Attributes and References.

Table 6-2 Alternate Value Encoding Pattern

Bounds	ValueType Encoding	Notes
[0..1]	Any(<type>)	An "optional" feature value with no elements is handled as follows: <ul style="list-style-type: none"> • the ref_get_value() raises Unset when the value is empty • the ref_unset_value() is used to set value to no elements
[1..1]	Any(<type>)	
others	Any(alias(seq(<type>, 0))) where the alias name is <typeName><CollectionKind>	A "multi-valued" feature value with no elements is encoded as a zero length sequence.

6.2.1.3 The Value Member Encoding Pattern

The following pattern is used in the reflective versions of the add, modify and remove operations that operate on the individual members of a multi-valued Attribute or Reference. The pattern is simply to encode the member as an Any containing an instance of the feature's base type. For example:

```
Any(<type>)
```

6.2.1.4 The Link Encoding Pattern

Some of the operations in the RefAssociation interface use the "generic" Link type to pass link values; see Section 6.3.2, "Data Types," on page 6-30. While the Link type uses RefObject rather than Any, a pattern is still required to describe the encoding.

The "generic" Link type is declared as a sequence of RefObject values with an upper bound of 2. The standard encoding of a link for a given Association is:

```
Link(<assocEnd1Type>, <assocEnd2Type>)
```

In other words, the sequence value contains precisely two elements, and the elements appear in the order of the corresponding AssociationEnds in the Association.

Reflective::RefBaseObject

abstract

The RefBaseObject interface is inherited by the other three reflective interfaces. It provides common operations for testing for object identity, returning an object's meta-object, and returning its "repository container" as required for implementing structural constraints such as the MOF's type closure rule and composition restrictions.

Supertypes

none (root object)

Operations

refMofId

The “refMofId” operation returns this object’s permanent unique identifier string.

<i>specific analog:</i>	none
<i>return type:</i>	string
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>exceptions:</i>	none

Every MOF object has a permanent, unique MOF identifier associated with it. This identifier is generated and bound to the object when it is created and cannot be changed for the lifetime of the object. The primary purpose of the MOF identifier is to serve as a label that can be compared to definitively establish an object’s identity.

A/n MOF implementation must ensure that no two distinct MOF objects within the extent of an outermost Package object ever have the same MOF identifier. This invariant must hold for the lifetime of the extent.

A group of outermost Package extents can only be safely federated if the respective implementations can ensure the above invariant applies across the entire federation. A federation of extents in which the invariant does not hold is not MOF compliant.

The MOF specification does not mandate a scheme for achieving this. Instead, the following approach is recommended:

1. Choose an appropriate scheme (or schemes) for allocating unique identifiers. This will depend on the nature of the federation.
2. Define a textual syntax for MOF identifier strings of the form:

```
<scheme-prefix> ":" <scheme-specific-part>
```

where <scheme-prefix> is either standardized elsewhere, or a vendor or user specific string that is unlikely to clash with other prefixes.

In the absence of a more appropriate identifier generation scheme, it is recommended that the following scheme based on the DCE UUID algorithm and textual encoding be used. The recommended DCE UUID-based identifier syntax is:

```
"DCE" ":" <printable-form-of-dce-uuid> [":" <decimal-digits>]
```

For example:

```
"DCE:d62207a2-011e-11ce-88b4-0800090b5d3e"
```

```
"DCE:d62207a2-011e-11ce-88b4-0800090b5d3e:1234"
```


The first case would be used when it is acceptable to generate a new DCE UUID for each MOF object. The second case might be used when the overheads of doing this are too large, or the required rate of UUID generation is too high. In this case, the UUID would denote an extent incarnation, and the suffix would be a local object sequence number for the extent incarnation does not repeat during the latter's lifetime.

refMetaObject

The “refMetaObject” operation returns the Model::ModelElement object that describes this object in its metamodel specification.

<i>specific analog:</i>	none
<i>return type:</i>	DesignatorType
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>exceptions:</i>	none

If the object's meta-object is unavailable, the return value may be a CORBA nil object reference.

refItself

The “refItself” operation tests whether this object and another RefBaseObject provided as an argument are the same CORBA object.

<i>specific analog:</i>	none
<i>return type:</i>	boolean
<i>isQuery:</i>	yes
<i>parameters:</i>	otherObject : in RefBaseObject
<i>exceptions:</i>	MofError (Invalid Object)

"Invalid Object" occurs if the "otherObject" is not a valid object, or if it is inaccessible.

refImmediatePackage

The “RefImmediatePackage” operation returns the RefPackage object for the Package that most immediately contains or aggregates this object.

<i>specific analog:</i>	none
<i>return type:</i>	RefPackage
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>exceptions:</i>	none

If this object has no containing or aggregating Package (i.e., it is the RefPackage object for an outermost Package), then the return value is a CORBA nil object reference. In complex cases where there is more than one immediate aggregating Package (see Section 4.6, “Extents,” on page 4-9 and Section 5.2.1, “Meta Object Type Overview,” on page 5-2), the return value may be any of them.

refOutermostPackage

The “refOutermostPackage” operation returns the RefPackage object for the Package that ultimately contains this object.

<i>specific analog:</i>	none
<i>return type:</i>	RefPackage
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>exceptions:</i>	none

If this object is the RefPackage object for an outermost Package, then the return value is this object.

refDelete

The “refDelete” operation destroys this object, including the objects it contains directly or transitively (see Section 5.3.2, “Lifecycle Semantics for the IDL Mapping,” on page 5-8 and Section 4.8, “Aggregation Semantics,” on page 4-17).

<i>specific analog:</i>	none
<i>return type:</i>	none
<i>parameters:</i>	none
<i>exceptions:</i>	MofError (Invalid Deletion)

The semantics of this operation depend on this RefBaseObject’s most derived type; see Section 5.2.1, “Meta Object Type Overview,” on page 5-2. Five sub-cases of RefBaseObject need to be considered here:

- outermost (i.e., non-nested, non-dependent) Package objects,
- nested or dependent Package objects,
- Association objects,
- Class proxy objects, and
- Instance objects.

Ordinary clients may only use “The “refDelete” operation destroys this object, including the objects it contains directly or transitively (see Section 5.3.2, “Lifecycle Semantics for the IDL Mapping,” on page 5-8 and Section 4.8, “Aggregation Semantics,” on page 4-17).” to delete instances of outermost Package objects and Instance objects.

- Deletion of an outermost Package causes all objects within its extent to be deleted; see Section 5.3.2.1, “Package object creation and deletion semantics,” on page 5-8.
- Deletion of an Instance object deletes it and its component closure; see Section 5.3.2.2, “Instance object lifecycle semantics,” on page 5-9.

“Invalid Deletion” occurs if an ordinary client invokes “refDelete” on a nested or dependent Package object, an Association object, or a Class proxy object.

As part of the deletion of an outermost Package, a Package object’s implementation may use the “refDelete” operation to delete nested or dependent Package objects, Association objects and Class proxy objects as well as Instance objects.

Interface

```

interface RefBaseObject {
  string ref_mof_id ();
  DesignatorType ref_meta_object ();
  boolean ref_itself (in RefBaseObject other_object);
  RefPackage ref_immediate_package ();
  RefPackage ref_outermost_package ();
  void ref_delete ()
    raises (MofError);
}; // end of RefBaseObject

```

6.2.2 *Reflective::RefObject*

Reflective::RefObject

abstract

The RefObject interface provides the meta-object description of an object that inherits from it, provides generic operations for testing for object identity and type membership, and a range of operations for accessing and updating the object in a model independent way.

The model assumed by the interface is that an object has structural features and operations. The model allows structural features to have single values or collection values. In the latter case, the collection values may have ordering or uniqueness semantics. There is provision for creation of new object instances, and for obtaining the set of objects that exist in a context.

Supertypes

RefBaseObject

Operations

refIsInstanceOf

This operation tests whether this RefObject is an instance of the Class described by the "someClass" meta-object. If the "considerSubtypes" argument is true, an object whose Class is a subclass of the Class described by "someClass" will be considered as an instance of the Class.

<i>specific analog:</i>	none
<i>return type:</i>	boolean
<i>isQuery:</i>	yes
<i>parameters:</i>	someClass : in DesignatorType considerSubtypes : in boolean
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind)

refCreateInstance

This operation creates a new instance of the Class for the RefObject's most derived interface. The operation can be called on a Class proxy object or an Instance object. The "args" list gives the initial values for the new Instance object's instance scoped, non-derived Attributes.

<i>specific analog:</i>	create_<class_name>(…); (see Section 5.8.9, “Class Create Template,” on page 5-58).
<i>return type:</i>	RefObject
<i>parameters:</i>	args : in ValueType (multiplicity: zero or more; ordered)
<i>exceptions:</i>	MofError (Overflow, Underflow, Duplicate, Composition Closure, Supertype Closure, Already Created, Abstract Class, Wrong Type, Wrong Number Parameters)

The members of the "args" list correspond 1-to-1 to the parameters for the specific create operation. They must be encoded as per Section 6.2.1.1, “The Standard Value Encoding Pattern,” on page 6-4. "Wrong Type" and "Wrong Number Parameters" when the "args" list has the wrong length or is incorrectly encoded.

"Abstract Class" occurs when “refCreateInstance” is called to create an instance of an "abstract Class. The remaining error conditions are directly equivalent to error conditions for the specific "create" operation.

refAllObjects

The “refAllObjects” operation returns the set of all Instances in the current extent whose type is given by this object’s Class. The operation can be called on a Class proxy object or an Instance object.

<i>specific analog:</i>	attribute all_of_type_<class_name>; attribute all_of_class_<class_name>; (See Section 5.8.6, “Class Template,” on page 5-54).
<i>return type:</i>	RefObject (multiplicity zero or more; unique; unordered)
<i>isQuery:</i>	yes
<i>parameters:</i>	includeSubtypes : in boolean
<i>exceptions:</i>	none

If "includeSubtypes" is true, the Instance objects for any subClasses of the M2 level Class are also included in the result set. This case is equivalent to the specific "all_of_type_<class_name>".

If the M2 level Class has “isAbstract” set to true, the result of

ref_all_objects(false)

is an empty set.

refValue

The “refValue” operation fetches the current value of the Attribute or Reference denoted by the "feature" argument. If this object is a Class proxy, only classifier scoped Attributes can be fetched.

<i>specific analog:</i>	<reference_name>(); (see Section 5.8.13, “Operation Template,” on page 5-87). <attribute_name>(); (see Section 5.8.11, “Attribute Template,” on page 5-68).
<i>return type:</i>	ValueType
<i>isQuery:</i>	yes
<i>parameters:</i>	feature : in DesignatorType
<i>exceptions:</i>	NotSet, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Underflow)

The result for the “refValue” operation is encoded as per Section 6.2.1.2, “The Alternate Value Encoding Pattern,” on page 6-4.

"NotSet" occurs when the feature's multiplicity is [0..1] and its value is unset; i.e. an empty collection. This should not occur with other multiplicities.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," and "Wrong Scope" all occur in cases where the "feature" argument does not denote an Attribute or Reference accessible from this object.

"Underflow" occurs when the feature is a Reference with multiplicity is [1..1] and its value has not been initialized. This should not occur for an Attribute or with other multiplicities.

refSetValue

The "refSetValue" operation assigns a new value to an Attribute or Reference for an object. The assigned value must be a single value or a collection value depending on the feature's multiplicity.

<i>specific analog:</i>	set_<reference_name>(newValue); (see Section 5.8.12, "Reference Template," on page 5-78). set_<attribute_name>(newValue); (see Section 5.8.11, "Attribute Template," on page 5-68).
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType newValue : in ValueType
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Underflow, Overflow, Duplicate, Reference Closure, Composition Closure, Composition Cycle, Invalid Object, Nil Object, Inaccessible Object, Wrong Type)

The "newValue" parameter must be encoded as per Section 6.2.1.2, "The Alternate Value Encoding Pattern," on page 6-4. "Wrong Type" occurs when this parameter is incorrectly encoded.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," and "Not Changeable" all occur in situations where the "feature" parameter does not denote a changeable Attribute or Reference that is accessible from this object.

The remaining error conditions are directly equivalent to error conditions for the "set_<feature_name>" operation.

refUnsetValue

The "refUnsetValue" operation resets an optional Attribute or Reference to contain no elements. This operation can only be used when the feature's multiplicity is [0..1].

<i>specific analog:</i>	unset_<reference_name>(); (see Section 5.8.12, "Reference Template," on page 5-78) unset_<attribute_name>(); (see Section 5.8.11, "Attribute Template," on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Underflow)

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur in situations where the "feature" parameter does not denote an Attribute or Reference for which "unset_<feature_name>" is allowed.

"Underflow" occurs in the same situation as for the "unset_<feature_name>" operation. For example, when "feature" is a Reference whose exposed Association End has a non-zero lower bound.

refAddValue

The "refAddValue" operation adds a new element to the current value of an Attribute or Reference with multiplicity that allows multiple values. If the Attribute or Reference is ordered, the new element is added at the end of the current value.

<i>specific analog:</i>	add_<reference_name>(newElement); (see Section 5.8.12, "Reference Template," on page 5-78) add_<attribute_name>(newElement); (see Section 5.8.11, "Attribute Template," on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType newElement : in ValueType
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "newElement" parameter should contain a single value of the feature's base type. "Wrong Type" occurs when it does not.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when the "feature" parameter does not designate a Reference or Attribute for which the "add_<feature_name>" operation is allowed.

The remaining error conditions are directly equivalent to error conditions for the "add_<feature_name>" operation.

refAddValueBefore

The "refAddValueBefore" operation is similar to "refAddValue" except that the caller specifies an existing element before which the new element is to be added. This operation can only be used for Attributes and References that are multi-valued and ordered. If the feature is non-unique (and therefore an Attribute), the insertion is made before the first element that matches, starting from the beginning of the collection.

<i>specific analog:</i>	add_<ref_name>_before(newElement, beforeElement); (see Section 5.8.12, "Reference Template," on page 5-78) add_<attr_name>_before(newElement, beforeElement); (see Section 5.8.11, "Attribute Template," on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType newElement : in ValueType beforeElement : in ValueType
<i>exceptions:</i>	NotFound, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "newElement" and "beforeElement" parameters should each contain a single value of the feature's base type. "Wrong Type" occurs when it does not.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when the "feature" parameter does not designate a Reference or Attribute for which the "add_<feature_name>_before" operation is allowed.

The remaining error conditions are directly equivalent to error conditions for the "add_<feature_name>_before" operation.

refAddValueAt

The “refAddValueAt” operation is similar to “refAddValueBefore” except that the caller explicitly gives the position of the insertion. The operation is only applicable to multi-valued ordered, non-unique Attributes.

<i>specific analog:</i>	add_<ref_name>_at(newElement, position); (see Section 5.8.12, “Reference Template,” on page 5-78) add_<attr_name>_at(newElement, position); (see Section 5.8.11, “Attribute Template,” on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType newElement : in ValueType position : in unsigned long
<i>exceptions:</i>	BadPosition, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "newElement" parameter should contain a single value of the Attribute’s base type. "Wrong Type" occurs if it is not.

The "position" parameter is interpreted the same way as for the corresponding specific operation. "Bad Position" occurs if the position parameter’s value is out of range, as defined for the "add_<feature_name>_at" operation (i.e., if it is greater than the size of the collection before the operation is invoke).

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when "feature" does not designate an Attribute for which the "add_<feature_name>_at" operation is allowed.

The remaining error conditions are directly equivalent to error conditions for the specific "add_<feature_name>_at" operation.

refModifyValue

The “refModifyValue” operation replaces one element of a multi-valued Attribute or Reference with a new value. If the feature is an ordered and non-unique (and therefore an Attribute), the element modified is the first one that matches, starting from the beginning of the collection.

<i>specific analog:</i>	modify_<ref_name>(oldElement, newElement); (see Section 5.8.12, "Reference Template," on page 5-78) modify_<attr_name>(oldElement, newElement); (see Section 5.8.11, "Attribute Template," on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType oldElement : in ValueType newElement : in ValueType
<i>exceptions:</i>	NotFound, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Underflow, Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "newElement" and "oldElement" parameters should contain a single value of the feature's base type. "Wrong Type" occurs if it is not.

The "oldElement" parameter should be an existing element of the collection being updated. "Not Found" occurs if it is not.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when the "feature" parameter does not designate a Reference or Attribute that supports the "modify_<feature_name>" operation.

The remaining error conditions are directly equivalent to error conditions for the "modify_<feature_name>" operation.

refModifyValueAt

The "refModifyValueAt" operation is similar to the "refModifyValue" operation, except that the element to be modified is specified by position. The operation is only applicable to multi-valued, ordered, non-unique Attributes.

<i>specific analog:</i>	modify_<ref_name>_at(newElement, position); (see Section 5.8.12, "Reference Template," on page 5-78) modify_<attr_name>_at(newElement, position); (see Section 5.8.11, "Attribute Template," on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType newElement : in ValueType position : in unsigned long
<i>exceptions:</i>	BadPosition, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Underflow, Overflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "newElement" parameter should contain a single value of the Attribute's base type. "Wrong Type" occurs if it is not.

The "position" parameter is interpreted in the same way as for the corresponding specific operation. "Bad Position" occurs if the position parameter's value is out of range, as defined for the "modify_<feature_name>_at" operation (i.e., if it is greater than or equal to the size of the collection).

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when "feature" does not designate an Attribute for which the "modify_<feature_name>_at" operation is allowed.

The remaining error conditions are directly equivalent to error conditions for the specific "modify_<feature_name>_at" operation.

refRemoveValue

The "refRemoveValue" operation removes an element of a multi-valued Attribute or Reference. The operation is only applicable when the upper bound is not equal to the lower bound. When the feature is ordered and non-unique (and therefore an Attribute) the element removed is the first one in the collection that matches, starting from the beginning of the collection.

<i>specific analog:</i>	remove_<reference_name>(oldElement); (see Section 5.8.12, "Reference Template," on page 5-78) remove_<attribute_name>(oldElement); (see Section 5.8.11, "Attribute Template," on page 5-68)
<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType oldElement : in ValueType
<i>exceptions:</i>	NotFound, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Underflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "oldElement" parameter should contain a single value of the Attribute's base type. "Wrong Type" occurs if it is not.

"Not Found" occurs if the value in the "oldElement" parameter is not a member of the collection.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when "feature" does not designate an Attribute or Reference for which the "remove_<feature_name>" operation is allowed.

The remaining error conditions are directly equivalent to error conditions for the specific "remove_<feature_name>" operation.

refRemoveValueAt

The "refRemoveValueAt" operation is similar to the "refRemoveValue" operation except that the element to be modified is specified by position. Furthermore, the operation is only applicable to ordered, non-unique Attributes.

<i>specific analog:</i>	remove_<reference_name>_at(position); (see Section 5.8.12, "Reference Template," on page 5-78) remove_<attribute_name>_at(position); (see Section 5.8.11, "Attribute Template," on page 5-68)
-------------------------	--

<i>return type:</i>	none
<i>parameters:</i>	feature : in DesignatorType position : in unsigned long
<i>exceptions:</i>	BadPosition, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Not Changeable, Wrong Multiplicity, Underflow, Duplicate, Invalid Object, Nil Object, Inaccessible Object, Reference Closure, Composition Closure, Composition Cycle, Wrong Type)

The "position" parameter is interpreted in the same way as for the corresponding specific operation. "Bad Position" occurs if the position parameter's value is out of range, as defined for the "remove_<feature_name>_at" operation (i.e., if it is greater than or equal to the size of the collection before the operation is called).

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," "Wrong Scope," "Not Changeable," and "Wrong Multiplicity" all occur when "feature" does not designate an Attribute for which the "remove_<feature_name>_at" operation is allowed.

The remaining error conditions are directly equivalent to error conditions for the specific "remove_<feature_name>_at" operation.

refImmediateComposite

The "refImmediateComposite" operation returns the "immediate composite" object for this Instance as specified below.

<i>specific analog:</i>	none
<i>return type:</i>	RefObject
<i>isQuery:</i>	yes
<i>exceptions:</i>	none

The immediate composite object C returned by this operation is an Instance object such that:

- C is related to this object via a relation R defined by an Attribute or Association,
- the aggregation semantics of the relation R are "composite", and
- this object fills the role of "component" in its relationship with C.

If the immediate object C does not exist, or if "this" object is a Class proxy object rather than an Instance object, a CORBA nil object reference is returned.

Note – If the composite relationship R corresponds to a "classifier-level" scoped M2-level Attribute, the immediate composite object C will be the Class Proxy object that holds the Attribute value.

refOutermostComposite

The “refOutermostComposite” operation returns the "outermost composite" for this object as defined below.

<i>specific analog:</i>	none
<i>return type:</i>	RefObject
<i>isQuery:</i>	yes
<i>exceptions:</i>	none

The outermost composite object C returned by this operation is an Instance object such that:

- There is a chain of *zero or more* immediate composite relationships (as described for “The “refImmediateComposite” operation returns the "immediate composite" object for this Instance as specified below.” above) connecting "this" object to C, and
- C does not have an immediate composite.

The above definition is such that if "this" object is not a component of any other object, it will be returned.

If "this" object is a Class proxy object, a CORBA nil object reference is returned.

Note – As with “refImmediateComposite” if the last composite relationship in the chain corresponds to a "classifier-level" scoped M2 level Attribute, the outermost composite object C will be the Class Proxy object that holds the Attribute value.

refInvokeOperation

The “refInvokeOperation” operation invokes a metamodel defined Operation on the Instance or Class proxy object with the arguments supplied.

<i>specific analog:</i>	none
<i>return type:</i>	ValueType (multiplicity: zero or more; ordered; not unique)
<i>parameters:</i>	requestedOperation : in DesignatorType args : inout ValueType (multiplicity: zero or more; ordered; non-unique)
<i>exceptions:</i>	OtherException, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Public, Wrong Scope, Overflow, Underflow, Duplicate, Wrong Number Parameters, Wrong Type)

The "args" parameter is used to pass the values of all of the Operation's Parameters which have directions "in," "out," or "inout" but not the "return" Parameter. There must be a distinct parameter value (real or dummy) in the "args" list for every "in," "out," and "inout" Parameter. "Wrong Number Parameters" occurs if this is not so.

The parameter values in "args" must appear in the order of the Operation's "in," "out," and "inout" Parameters as defined in the metamodel.

The "args" member values provided by the caller for "in" and "inout" Parameter positions must be encoded depending on the Parameter's type and multiplicity as per the Section 6.2.1.1, "The Standard Value Encoding Pattern," on page 6-4. "Wrong Type" occurs if any of these values have the wrong type for the corresponding Parameter. "Underflow," "Overflow," or "Duplicate" occur when one of the supplied values does not fit the multiplicity specified by the corresponding Parameter.

The caller must provide a dummy "args" member value in each "out" Parameter position. This value may be any legal CORBA Any value.

The "args" member values passed back to the caller for "out" and "inout" Parameter positions are likewise encoded depending on the Parameter's type and multiplicity as per the Section 6.2.1.1, "The Standard Value Encoding Pattern," on page 6-4. Note that the values passed back to the caller the "in" Parameter positions of the "args" list are dummies whose content is undefined.

If the Operation defines a result (i.e., a Parameter with direction "return"), the result for a "The "refInvokeOperation" operation invokes a metamodel defined Operation on the Instance or Class proxy object with the arguments supplied." call gives the result value. This is encoded depending on the "return" Parameter's type and multiplicity as per the Section 6.2.1.1, "The Standard Value Encoding Pattern," on page 6-4. When the Operation does not define a result, the result of a "The "refInvokeOperation" operation invokes a metamodel defined Operation on the Instance or Class proxy object with the arguments supplied." call is a dummy value whose content is undefined.

Note – In the cases above where dummy values are used, it is recommended that "light weight" Any values are used. (We would recommend the use of an Any value whose type kind is tk_null. However, there is currently some question as to whether the CDR standard defines an encoding for this value.)

"OtherException" occurs when a "refInvokeOperation" invocation needs to signal an Operation specific Exception. The "exception_designator" field of "OtherException" will denote the Exception raised, and the "exception_args" list will give the values for any Exception fields. The "exception_args" list will have one member value for each Parameter of the Exception in the order defined by the meta-model. The member values will be encoded depending on the corresponding Exception Parameter's type and multiplicity as per the Section 6.2.1.1, "The Standard Value Encoding Pattern," on page 6-4.

"Invalid Designator," "Wrong Designator Kind," "Unknown Designator," "Not Public," and "Wrong Scope" all occur when "requestedOperation" does not designate an Operation that can be invoked using this object.

Interface

```
interface RefObject : RefBaseObject {
    boolean ref_is_instance_of (in DesignatorType some_class,
                               in boolean consider_subtypes);
    RefObject ref_create_instance (in ValueTypeList args)
        raises (MofError);
    RefObjectSet ref_all_objects (in boolean include_subtypes);
    void ref_set_value (in DesignatorType feature,
                      in ValueType new_value)
        raises (MofError);
    ValueType ref_value (in DesignatorType feature)
        raises (NotSet, MofError);
    void ref_unset_value ()
        raises (MofError);
    void ref_add_value (in DesignatorType feature,
                      in ValueType new_element)
        raises (MofError);
    void ref_add_value_before (in DesignatorType feature,
                              in ValueType new_element,
                              in ValueType before_element)
        raises (NotFound, MofError);
    void ref_add_value_at (in DesignatorType feature,
                          in ValueType new_element,
                          in unsigned long position)
        raises (BadPosition, MofError);
    void ref_modify_value (in DesignatorType feature,
                          in ValueType old_element,
                          in ValueType new_element)
        raises (NotFound, MofError);
    void ref_modify_value_at (in DesignatorType feature,
                              in ValueType new_element,
                              in unsigned long position)
        raises (BadPosition, MofError);
    void ref_remove_value (in DesignatorType feature,
                          in ValueType old_element)
        raises (NotFound, MofError);
    void ref_remove_value_at (in DesignatorType feature,
                              in unsigned long position)
        raises (BadPosition, MofError);
    RefObject ref_immediate_composite ();
    RefObject ref_outermost_composite ();
    ValueType ref_invoke_operation (
        in DesignatorType requested_operation,
        inout ValueTypeList args)
        raises (OtherException, MofError);
}; // end of interface RefObject
```


6.2.3 *Reflective::RefAssociation*

Reflective::RefAssociation

abstract

The RefAssociation interface provides the meta-object description of an association that inherits from it. It also provides generic operations querying and updating the links that belong to the association.

The model of association supported by this interface is of collection of two ended asymmetric links between objects. The links may be viewed as ordered on one or other of the ends, and there may be some form of cardinality constraints on either end.

The RefAssociation interface is designed to be used with associations that contain no duplicate links, though this is not an absolute requirement. There is no assumption that different association objects for a given association type are mutually aware. Links are modeled as having no object identity.

A data model that required "heavy weight" links with object identity (e.g., so that attributes could be attached to them) would need to represent them as RefObject instances. The RefAssociation interface could be used to manage light weight links between heavy weight link objects and the objects they connect. Similar techniques could be used to represent N-ary associations. However, in both cases better performance would be achieved using a purpose built reflective layer.

Supertypes

RefBaseObject

Operations

refAllLinks

The "refAllLinks" operation returns all links in the link set for this Association object.

<i>specific analog:</i>	all_links(); (see Section 5.8.10, "Association Template," on page 5-59)
<i>return type:</i>	Link (multiplicity zero or more, unordered, unique)
<i>isQuery:</i>	yes
<i>parameters:</i>	none
<i>exceptions:</i>	none

This operation returns the current link set for the current Association extent as defined for the specific version of this operation. The links are encoded as per the Section 6.2.1.4, "The Link Encoding Pattern," on page 6-5.

refLinkExists

The “refLinkExists” operation returns true if and only if the supplied link is a member of the link set for this Association object.

<i>specific analog:</i>	link_exists(someLink); (see Section 5.8.10, “Association Template,” on page 5-59)
<i>return type:</i>	boolean
<i>isQuery:</i>	yes
<i>parameters:</i>	someLink : in Link
<i>exceptions:</i>	MofError(WrongType)

The "someLink" parameter should be encoded as per the Section 6.2.1.4, “The Link Encoding Pattern,” on page 6-5. "Wrong Type" occurs if the link encoding is not correct.

refQuery

The “refQuery” operation returns a list containing all Instance objects that are linked to the supplied "queryObject" by links in the extent of this Association object, where the links all have the "queryObject" at the "queryEnd."

<i>specific analog:</i>	<endName> (queryObject); (see Section 5.8.10, “Association Template,” on page 5-59)
<i>return type:</i>	RefObject (Multiplicity zero or more; ordered; unique)
<i>isQuery:</i>	yes
<i>parameters:</i>	queryEnd : in DesignatorType queryObject : in RefObject
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Wrong Type, Invalid Object, Nil Object, Inaccessible Object)

The "queryEnd" parameter must designate an AssociationEnd for this Association object. "Invalid Designator," "Wrong Designator Kind," and "Unknown Designator" occur in cases where this is not so.

The "queryObject" parameter must be an Instance object whose type is compatible with the type of the "queryEnd" of the Association. "Wrong Type" is raised if it the parameter has the wrong type.

"Invalid Object," "Nil Object," or "Inaccessible Object" is raised if the "queryObject" parameter it is a non-existent, nil or inaccessible Instance object.

While the result of this operation is declared as a ordered set of links, the ordering only has meaning if the other AssociationEnd (i.e., not the "queryEnd") is defined ordered.

refAddLink

The “refAddLink” operation adds "newLink" into the set of links in the extent of this Association object. If one or other of the Association’s Ends is ordered, the link is inserted after the last link with respect to that ordering.

<i>specific analog:</i>	add(newLink[0], newLink[1]); (see Section 5.8.10, “Association Template,” on page 5-59)
<i>return type:</i>	none
<i>parameters:</i>	newLink : in Link
<i>exceptions:</i>	MofError (Not Changeable, Overflow, Duplicate, Reference Closure, Composition Closure, Composition Cycle, Wrong Type, Invalid Object, Nil Object, Inaccessible Object)

The "newLink" parameter should be encoded as per the Section 6.2.1.4, “The Link Encoding Pattern,” on page 6-5. "Wrong Type" occurs if the link encoding is not correct.

Both RefObject members of the "newLink" parameter should be valid Instance objects. "Invalid Object," "Nil Object," or "Inaccessible Object" is raised if either one is a non-existent, nil or inaccessible Instance object.

"Not Changeable" occurs if this operation is invoked on an Association that has “isChangeable” set to false on either Association End.

"Overflow," "Duplicate," "Reference Closure," "Composition Closure," and "Composition Cycle" are directly equivalent to error conditions for the corresponding specific "add" operation.

refAddLinkBefore

The “refAddLinkBefore” operation adds "newLink" into the link set of an ordered Association object. The link insertion point is immediately before the link whose "positionEnd" matches the "before" Instance.

<i>specific analog:</i>	add_before_<endName>(newLink[0], newLink[1], before); (see Section 5.8.10, “Association Template,” on page 5-59)
-------------------------	---

<i>return type:</i>	none
<i>parameters:</i>	newLink : in Link positionEnd : in DesignatorType before : in RefObject
<i>exceptions:</i>	NotFound, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Changeable, Not Navigable, Overflow, Duplicate, Reference Closure, Composition Closure, Wrong Type, Invalid Object, Nil Object, Inaccessible Object)

The "newLink" parameter should be encoded as per the Section 6.2.1.4, "The Link Encoding Pattern," on page 6-5. "Wrong Type" occurs if the link's encoding is not correct.

The "positionEnd" parameter should denote an AssociationEnd of this object's Association. One of "Invalid Designator," "Wrong Designator Kind," or "Unknown Designator" occurs if this is not the case.

"Not Changeable" occurs if this operation is invoked on an Association that has "isChangeable" set to false on either Association End. "Not Navigable" occurs if the "positionEnd" AssociationEnd has "isNavigable" set to false.

The "before" parameter should be an Instance object that is type compatible with the type of the AssociationEnd denoted by "positionEnd." "Wrong Type" occurs if this is not the case.

The remaining error conditions are directly equivalent to error conditions for the corresponding "add_before_<endName>" operation.

refModifyLink

The "refModifyLink" operation updates the "oldLink" in the Association object's link set, replacing the Instance object at "positionEnd" with "newObject."

<i>specific analog:</i>	modify_<endName>(oldLink[0], oldLink[1], newObject); (see Section 5.8.10, "Association Template," on page 5-59)
<i>return type:</i>	none
<i>parameters:</i>	oldLink : in Link positionEnd : in DesignatorType newObject : in RefObject
<i>exceptions:</i>	NotFound, MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator, Not Changeable, Underflow, Overflow, Duplicate, Reference Closure, Composition Closure, Wrong Type, Invalid Object, Nil Object, Inaccessible Object)

The "oldLink" parameter should be encoded as per the Section 6.2.1.4, "The Link Encoding Pattern," on page 6-5. "Wrong Type" occurs if the link's encoding is not correct.

The "positionEnd" parameter should denote an AssociationEnd of this object's Association. One of "Invalid Designator," "Wrong Designator Kind," or "Unknown Designator" occurs if this is not the case.

"Not Changeable" occurs if the "positionEnd" AssociationEnd that has "isChangeable" set to false. "Not Navigable" occurs if it has "isNavigable" set to false.

The "newObject" parameter should be an Instance object that is type compatible with the type of the AssociationEnd denoted by "positionEnd." "Wrong Type" occurs if this is not the case.

The remaining error conditions are directly equivalent to error conditions for the corresponding "modify_<endName>" operation. Note that any structural constraints notionally apply to the final state following the operation, and not to any intermediate states.

refRemoveLink

The "refRemoveLink" operation removes the "oldLink" from the association.

<i>specific analog:</i>	remove(oldLink[0], oldLink[1]); (see Section 5.8.10, "Association Template," on page 5-59)
<i>return type:</i>	none
<i>parameters:</i>	oldLink : in Link
<i>exceptions:</i>	NotFound, MofError (Not Changeable, Underflow, Wrong Type, Nil Object)

"Not Changeable" occurs if this operation is invoked on an Association that has "isChangeable" set to false for either AssociationEnd.

The "oldLink" parameter should be encoded as per the Section 6.2.1.4, "The Link Encoding Pattern," on page 6-5. "Wrong Type" occurs if the link's encoding is not correct.

"NotFound," "Nil Object," and "Underflow" are directly equivalent to error conditions for the corresponding specific "remove" operation. "Invalid Object" and "Inaccessible Object" cannot occur, as in the specific operation.

Interface

```
interface RefAssociation : RefBaseObject {
    LinkSet ref_all_links ();
    boolean ref_link_exists (in Link some_link)
        raises (MofError);
    RefObjectUList ref_query (in DesignatorType query_end,
        in RefObject query_object)
```

```

    raises (MofError);
void ref_add_link (in Link new_link)
    raises (MofError);
void ref_add_link_before (in Link new_link,
                        in DesignatorType position_end,
                        in RefObject before)
    raises (NotFound, MofError);
void ref_modify_link (in Link old_link,
                    in DesignatorType position_end,
                    in RefObject new_object)
    raises (NotFound, MofError);
void ref_remove_link (in Link old_link)
    raises (NotFound, MofError);
}; // end of interface RefAssociation

```

6.2.4 Reflective::RefPackage

Reflective::RefPackage

abstract

The RefPackage interface is an abstraction for accessing a collection of objects and their associations. The interface provides an operation to access the meta-object description for the package, and operations to access the package instance's class proxy objects (one for each Class) and its association objects.

Supertypes

RefBaseObject

Operations

refClassRef

The “refClassRef” operation returns the Class proxy object for a given Class.

<i>specific analog:</i>	readonly attribute <ClassName>_class_ref; (see Section 5.8.10, “Association Template,” on page 5-59)
<i>return type:</i>	RefObject
<i>isQuery:</i>	yes
<i>parameters:</i>	class : in DesignatorType
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator)

The "class" parameter should designate the M2 level Class whose Class proxy object is to be returned. "Invalid Designator," "Wrong Designator Kind," "Unknown Designator" occur in various situations where this is not the case.

refAssociationRef

The “refAssociationRef” operation returns an Association object for a given Association.

<i>specific analog:</i>	readonly attribute <AssociationName>_ref; (see Section 5.8.10, “Association Template,” on page 5-59)
<i>return type:</i>	RefAssociation
<i>isQuery:</i>	yes
<i>parameters:</i>	association : DesignatorType
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator)

The "association" parameter should designate the M2 level Association whose Association object is to be returned. "Invalid Designator," "Wrong Designator Kind," "Unknown Designator" occur in various situations where this is not the case.

refPackageRef

The “refPackageRef” operation returns a Package object for a nested or clustered Package.

<i>specific analog:</i>	readonly attribute <PackageName>_ref; (see Section 5.8.10, “Association Template,” on page 5-59)
<i>return type:</i>	RefPackage
<i>isQuery:</i>	yes
<i>parameters:</i>	package : DesignatorType
<i>exceptions:</i>	MofError (Invalid Designator, Wrong Designator Kind, Unknown Designator)

The "package" parameter should designate the M2 level Package whose Package object is to be returned. It must either be nested within the Package for this Package object, or imported with “isClustered” set to true. "Invalid Designator," "Wrong Designator Kind," "Unknown Designator" occur in the situations where this is not the case.

Interface

```
interface RefPackage : RefBaseObject {
```

```
  RefObject ref_class_ref (in DesignatorType type)  
    raises (MofError);
```

```
  RefAssociation ref_association_ref (  
    in DesignatorType association)  
    raises (MofError);
```

```
RefPackage ref_package_ref (in DesignatorType package)
raises (InvalidDesignator)

}; // end of interface RefPackage
```

6.3 *The CORBA IDL for the Reflective Interfaces*

This section describes the relevant excerpts of the CORBA IDL for the Reflective module.

6.3.1 *Introduction*

The Reflective module starts with forward declarations of the three object types RefObject, RefAssociation, and RefPackage.

```
module Reflective {
  interface RefBaseObject;

  interface RefObject;
  typedef sequence < RefObject > RefObjectUList;

  interface RefAssociation;

  interface RefPackage;
```

6.3.2 *Data Types*

Operations on the Reflective interfaces need to identify the elements (e.g., attributes, operations, roles, classes, etc.) that they apply to. Some exceptions have similar requirements. The type DesignatorType is used to denote uses of RefObject with this meaning.

```
typedef RefObject DesignatorType;
```

Values of attributes, operation parameters, and results etc. are passed using the CORBA "any" data type. The type ValueType is used to denote uses of "any" with this meaning. The encoding of values using the "any" type is model specific.

```
typedef any ValueType;
```

Links are expressed as bounded sequences of (two) RefObject values.

```
typedef sequence <RefObject, 2> Link;
typedef sequence <Link> LinkSet;
```


Glossary

This glossary defines the terms that are used to describe the Unified Modeling Language (UML) and the Meta Object Facility (MOF). In addition to UML and MOF specific terminology, it includes related terms from OMG standards and object-oriented analysis and design methods, as well as the domain of object repositories and meta data managers. Glossary entries are organized alphabetically and MOF specific entries are identified as '[MOF]'.

Notation Conventions

The entries in the glossary usually begin with a lowercase letter. An initial uppercase letter is used when a word is usually capitalized in standard practice. Acronyms are all capitalized, unless they traditionally appear in all lowercase.

When one or more words in a multi-word term is enclosed in brackets, it indicates that those words are optional when referring to the term. For example, *use case [class]* may be referred to as simply *use case*.

The following conventions are used in this glossary:

- Contrast: <term>
Refers to a term that has an opposed or substantively different meaning.
- See: <term>
Refers to a related term that has a similar, but not synonymous meaning.
- Synonym: <term>
Indicates that the term has the same meaning as another term, which is referenced.
- Acronym: <term>
Indicates that the term is an acronym. The reader is usually referred to the spelled-out term for the definition, unless the spelled-out term is rarely used.

Glossary Terms

abstract class	A class that cannot be directly instantiated. Contrast: <i>concrete class</i> .
abstraction	The essential characteristics of an entity that distinguish it from all other kinds of entities. An abstraction defines a boundary relative to the perspective of the viewer.
action	The specification of an executable statement that forms an abstraction of a computational procedure. An action typically results in a change in the state of the system, and can be realized by sending a message to an object or modifying a link or a value of an attribute.
action sequence	An expression that resolves to a sequence of actions.
action state	A state that represents the execution of an atomic action, typically the invocation of an operation.
activation	The execution of an action.
active class	A class whose instances are active objects. See: <i>active object</i> .
active object	An object that owns a thread and can initiate control activity. An instance of active class. See: <i>active class</i> , <i>thread</i> .
activity graph	A special case of a state machine that is used to model processes involving one or more classifiers. Contrast: <i>statechart diagram</i> .
actor [class]	A coherent set of roles that users of use cases play when interacting with these use cases. An actor has one role for each use case with which it communicates.
actual parameter	Synonym: <i>argument</i> .

aggregate [class]	A class that represents the “whole” in an aggregation (whole-part) relationship. See: <i>aggregation</i> .
aggregation	A special form of association that specifies a whole-part relationship between the aggregate (whole) and a component part. See: <i>composition</i> .
analysis	The part of the software development process whose primary purpose is to formulate a model of the problem domain. Analysis focuses what to do, design focuses on how to do it. Contrast: <i>design</i> .
analysis time	Refers to something that occurs during an analysis phase of the software development process. See: <i>design time, modeling time</i> .
architecture	The organizational structure and associated behavior of a system. An architecture can be recursively decomposed into parts that interact through interfaces, relationships that connect parts, and constraints for assembling parts. Parts that interact through interfaces include classes, components and subsystems.
argument	A binding for a parameter that resolves to a run-time instance. Synonym: <i>actual parameter</i> . Contrast: <i>parameter</i> .
artifact	A piece of information that is used or produced by a software development process. An artifact can be a model, a description, or software. Synonym: <i>product</i> .
association	The semantic relationship between two or more classifiers that specifies connections among their instances.
association class	A model element that has both association and class properties. An association class can be seen as an association that also has class properties, or as a class that also has association properties.
association end	The endpoint of an association, which connects the association to a classifier.

attribute	A feature within a classifier that describes a range of values that instances of the classifier may hold.
behavior	The observable effects of an operation or event, including its results.
behavioral feature	A dynamic feature of a model element, such as an operation or method.
behavioral model aspect	A model aspect that emphasizes the behavior of the instances in a system, including their methods, collaborations, and state histories.
binary association	An association between two classes. A special case of an n-ary association.
binding	The creation of a model element from a template by supplying arguments for the parameters of the template.
boolean	An enumeration whose values are true and false.
boolean expression	An expression that evaluates to a boolean value.
cardinality	The number of elements in a set. Contrast: <i>multiplicity</i> .
child	In a generalization relationship, the specialization of another element, the parent. See: <i>subclass</i> , <i>subtype</i> . Contrast: <i>parent</i> .
call	An action state that invokes an operation on a classifier.
class	A description of a set of objects that share the same attributes, operations, methods, relationships, and semantics. A class may use a set of interfaces to specify collections of operations it provides to its environment. See: <i>interface</i> .
classifier	A mechanism that describes behavioral and structural features. Classifiers include interfaces, classes, datatypes, and components.

classification	The assignment of an object to a classifier. See <i>dynamic classification</i> , <i>multiple classification</i> and <i>static classification</i> .
class diagram	A diagram that shows a collection of declarative (static) model elements, such as classes, types, and their contents and relationships.
client	A classifier that requests a service from another classifier. Contrast: <i>supplier</i> .
collaboration	The specification of how an operation or classifier, such as a use case, is realized by a set of classifiers and associations playing specific roles used in a specific way. The collaboration defines an interaction. See: <i>interaction</i> .
collaboration diagram	A diagram that shows interactions organized around the structure of a model, using either classifiers and associations or instances and links. Unlike a sequence diagram, a collaboration diagram shows the relationships among the instances. Sequence diagrams and collaboration diagrams express similar information, but show it in different ways. See: <i>sequence diagram</i> .
comment	An annotation attached to an element or a collection of elements. A note has no semantics. Contrast: <i>constraint</i> .
compile time	Refers to something that occurs during the compilation of a software module. See: <i>modeling time</i> , <i>run time</i> .
component	A physical, replaceable part of a system that packages implementation and provides the realization of a set of interfaces. A component represents a physical piece of implementation of a system, including software code (source, binary or executable) or equivalents such as scripts or command files.

component diagram	A diagram that shows the organizations and dependencies among components.
composite [class]	A class that is related to one or more classes by a composition relationship. See: <i>composition</i> .
composite aggregation	Synonym: <i>composition</i> .
composite state	A state that consists of either concurrent (orthogonal) substates or sequential (disjoint) substates. See: <i>substate</i> .
composition	A form of aggregation association with strong ownership and coincident lifetime as part of the whole. Parts with non-fixed multiplicity may be created after the composite itself, but once created they live and die with it (i.e., they share lifetimes). Such parts can also be explicitly removed before the death of the composite. Composition may be recursive. Synonym: <i>composite aggregation</i> .
concrete class	A class that can be directly instantiated. Contrast: <i>abstract class</i> .
concurrency	The occurrence of two or more activities during the same time interval. Concurrency can be achieved by interleaving or simultaneously executing two or more threads. See: <i>thread</i> .
concurrent substate	A substate that can be held simultaneously with other substates contained in the same composite state. See: <i>composite state</i> . Contrast: <i>disjoint substate</i> .
constraint	A semantic condition or restriction. Certain constraints are predefined in the UML, others may be user defined. Constraints are one of three extensibility mechanisms in UML. See: <i>tagged value</i> , <i>stereotype</i> .
container	1. An instance that exists to contain other instances, and that provides operations to access or iterate over its contents. (for example, arrays, lists, sets). 2. A component that exists to contain other components.

containment hierarchy	A namespace hierarchy consisting of model elements, and the containment relationships that exist between them. A containment hierarchy forms a graph.
context	A view of a set of related modeling elements for a particular purpose, such as specifying an operation.
datatype	A descriptor of a set of values that lack identity and whose operations do not have side effects. Datatypes include primitive pre-defined types and user-definable types. Pre-defined types include numbers, string and time. User-definable types include enumerations.
defining model [MOF]	The model on which a repository is based. Any number of repositories can have the same defining model.
delegation	The ability of an object to issue a message to another object in response to a message. Delegation can be used as an alternative to inheritance. Contrast: <i>inheritance</i> .
dependency	A relationship between two modeling elements, in which a change to one modeling element (the independent element) will affect the other modeling element (the dependent element).
deployment diagram	A diagram that shows the configuration of run-time processing nodes and the components, processes, and objects that live on them. Components represent run-time manifestations of code units. See: <i>component diagrams</i> .
derived element	A model element that can be computed from another element, but that is shown for clarity or that is included for design purposes even though it adds no semantic information.

design	The part of the software development process whose primary purpose is to decide how the system will be implemented. During design strategic and tactical decisions are made to meet the required functional and quality requirements of a system.
design time	Refers to something that occurs during a design phase of the software development process. See: <i>modeling time</i> . Contrast: <i>analysis time</i> .
development process	A set of partially ordered steps performed for a given purpose during software development, such as constructing models or implementing models.
diagram	A graphical presentation of a collection of model elements, most often rendered as a connected graph of arcs (relationships) and vertices (other model elements). UML supports the following diagrams: class diagram, object diagram, use case diagram, sequence diagram, collaboration diagram, state diagram, activity diagram, component diagram, and deployment diagram.
disjoint substate	A substate that cannot be held simultaneously with other substates contained in the same composite state. See: composite state. Contrast: <i>concurrent substate</i> .
distribution unit	A set of objects or components that are allocated to a process or a processor as a group. A distribution unit can be represented by a run-time composite or an aggregate.
domain	An area of knowledge or activity characterized by a set of concepts and terminology understood by practitioners in that area.
dynamic classification	A semantic variation of generalization in which an object may change its classifier. Contrast: <i>static classification</i> .
element	An atomic constituent of a model.

entry action	An action executed upon entering a state in a state machine regardless of the transition taken to reach that state.
enumeration	A list of named values used as the range of a particular attribute type. For example, RGBColor = {red, green, blue}. Boolean is a predefined enumeration with values from the set {false, true}.
event	The specification of a significant occurrence that has a location in time and space. In the context of state diagrams, an event is an occurrence that can trigger a transition.
exit action	An action executed upon exiting a state in a state machine regardless of the transition taken to exit that state.
export	In the context of packages, to make an element visible outside its enclosing namespace. See: <i>visibility</i> . Contrast: <i>export</i> [OMA], <i>import</i> .
expression	A string that evaluates to a value of a particular type. For example, the expression “(7 + 5 * 3)” evaluates to a value of type number.
extend	A relationship from an extension use case to a base use case, specifying how the behavior defined for the extension use case augments (subject to conditions specified in the extension) the behavior defined for the base use case. The behavior is inserted at the location defined by the extension point in the base use case. The base use case does not depend on performing the behavior of the extension use case. See <i>extension point</i> , <i>include</i> .
facade	A stereotyped package containing only references to model elements owned by another package. It is used to provide a ‘public view’ of some of the contents of a package.
feature	A property, like operation or attribute, which is encapsulated within a classifier, such as an interface, a class, or a datatype.

final state	A special kind of state signifying that the enclosing composite state or the entire state machine is completed.
fire	To execute a state transition. See: <i>transition</i> .
focus of control	A symbol on a sequence diagram that shows the period of time during which an object is performing an action, either directly or through a subordinate procedure.
formal parameter	Synonym: <i>parameter</i> .
framework	<ol style="list-style-type: none"> 1. A stereotyped package consisting mainly of patterns. See: <i>pattern</i>. 2. An architectural pattern that provides an extensible template for applications within a specific domain.
generalizable element	A model element that may participate in a generalization relationship. See: <i>generalization</i> .
generalization	A taxonomic relationship between a more general element and a more specific element. The more specific element is fully consistent with the more general element and contains additional information. An instance of the more specific element may be used where the more general element is allowed. See: <i>inheritance</i> .
guard condition	A condition that must be satisfied in order to enable an associated transition to fire.
implementation	A definition of how something is constructed or computed. For example, a class is an implementation of a type, a method is an implementation of an operation.
implementation inheritance	The inheritance of the implementation of a more specific element. Includes inheritance of the interface. Contrast: <i>interface inheritance</i> .

import	In the context of packages, a dependency that shows the packages whose classes may be referenced within a given package (including packages recursively embedded within it). Contrast: <i>export</i> .
include	A relationship from a base use case to an inclusion use case, specifying how the behavior for the base use case contains the behavior of the inclusion use case. The behavior is included at the location which is defined in the base use case. The base use case depends on performing the behavior of the inclusion use case, but not on its structure (i.e., attributes or operations). See <i>extend</i> .
inheritance	The mechanism by which more specific elements incorporate structure and behavior of more general elements related by behavior. See <i>generalization</i> .
instance	An entity to which a set of operations can be applied and which has a state that stores the effects of the operations. See: <i>object</i> .
interaction	A specification of how stimuli are sent between instances to perform a specific task. The interaction is defined in the context of a collaboration. See <i>collaboration</i> .
interaction diagram	A generic term that applies to several types of diagrams that emphasize object interactions. These include collaboration diagrams and sequence diagrams.
interface	A named set of operations that characterize the behavior of an element.
interface inheritance	The inheritance of the interface of a more specific element. Does not include inheritance of the implementation. Contrast: <i>implementation inheritance</i> .
internal transition	A transition signifying a response to an event without changing the state of an object.

layer	The organization of classifiers or packages at the same level of abstraction. A layer represents a horizontal slice through an architecture, whereas a partition represents a vertical slice. Contrast: <i>partition</i> .
link	A semantic connection among a tuple of objects. An instance of an association. See: <i>association</i> .
link end	An instance of an association end. See: <i>association end</i> .
message	A specification of the conveyance of information from one instance to another, with the expectation that activity will ensue. A message may specify the raising of a signal or the call of an operation.
metaclass	A class whose instances are classes. Metaclasses are typically used to construct metamodels.
meta-metamodel	A model that defines the language for expressing a metamodel. The relationship between a meta-metamodel and a metamodel is analogous to the relationship between a metamodel and a model.
metamodel	A model that defines the language for expressing a model.
metaobject	A generic term for all metaentities in a metamodeling language. For example, metatypes, metaclasses, metaattributes, and metaassociations.
method	The implementation of an operation. It specifies the algorithm or procedure associated with an operation.
model	An abstraction of a physical system, with a certain purpose.. See: <i>physical system</i> .
[MOF]	Usage note: In the context of the MOF specification, which describes a meta-metamodel, for brevity the meta-metamodel is frequently to as simply the model.

model aspect	A dimension of modeling that emphasizes particular qualities of the metamodel. For example, the structural model aspect emphasizes the structural qualities of the metamodel.
model elaboration	The process of generating a repository type from a published model. Includes the generation of interfaces and implementations which allows repositories to be instantiated and populated based on, and in compliance with, the model elaborated.
model element	An element that is an abstraction drawn from the system being modeled. Contrast: <i>view element</i> .
[MOF]	In the MOF specification model elements are considered to be metaobjects.
modeling time	Refers to something that occurs during a modeling phase of the software development process. It includes analysis time and design time. Usage note: When discussing object systems, it is often important to distinguish between modeling-time and run-time concerns. See: <i>analysis time, design time</i> . Contrast: <i>run time</i> .
module	A software unit of storage and manipulation. Modules include source code modules, binary code modules, and executable code modules. See: <i>component</i> .
multiple classification	A semantic variation of generalization in which an object may belong directly to more than one classifier. See: <i>static classification, dynamic classification</i> .
multiple inheritance	A semantic variation of generalization in which a type may have more than one supertype. Contrast: <i>single inheritance</i> .

multiplicity	A specification of the range of allowable cardinalities that a set may assume. Multiplicity specifications may be given for roles within associations, parts within composites, repetitions, and other purposes. Essentially a multiplicity is a (possibly infinite) subset of the non-negative integers. Contrast: <i>cardinality</i> .
multi-valued [MOF]	A model element with multiplicity defined whose Multiplicity Type:: upper attribute is set to a number greater than one. The term multi-valued does not pertain to the number of values held by an attribute, parameter, etc. at any point in time. Contrast: <i>single-valued</i> .
n-ary association	An association among three or more classes. Each instance of the association is an n-tuple of values from the respective classes. Contrast: <i>binary association</i> .
name	A string used to identify a model element.
namespace	A part of the model in which the names may be defined and used. Within a namespace, each name has a unique meaning. See: <i>name</i> .
node	A node is classifier that represents a run-time computational resource, which generally has at least a memory and often processing capability. Run-time objects and components may reside on nodes.
object	An entity with a well-defined boundary and identity that encapsulates state and behavior. State is represented by attributes and relationships, behavior is represented by operations, methods, and state machines. An object is an instance of a class. See: <i>class, instance</i> .
object diagram	A diagram that encompasses objects and their relationships at a point in time. An object diagram may be considered a special case of a class diagram or a collaboration diagram. See: <i>class diagram, collaboration diagram</i> .

object flow state	A state in an activity graph that represents the passing of an object from the output of actions in one state to the input of actions in another state.
object lifeline	A line in a sequence diagram that represents the existence of an object over a period of time. See: <i>sequence diagram</i> .
operation	A service that can be requested from an object to effect behavior. An operation has a signature, which may restrict the actual parameters that are possible.
package	A general purpose mechanism for organizing elements into groups. Packages may be nested within other packages.
parameter	The specification of a variable that can be changed, passed, or returned. A parameter may include a name, type, and direction. Parameters are used for operations, messages, and events. Synonyms: <i>formal parameter</i> . Contrast: <i>argument</i> .
parameterized element	The descriptor for a class with one or more unbound parameters. Synonym: <i>template</i> .
parent	In a generalization relationship, the generalization of another element, the child. See: <i>subclass</i> , <i>subtype</i> . Contrast: <i>child</i> .
participate	The connection of a model element to a relationship or to a reified relationship. For example, a class participates in an association, an actor participates in a use case.
partition	<ol style="list-style-type: none"> 1. activity graphs: A portion of an activity graphs that organizes the responsibilities for actions. See: <i>swimlane</i>. 2. architecture: A set of related classifiers or packages at the same level of abstraction or across layers in a layered architecture. A partition represents a vertical slice through an architecture, whereas a layer represents a horizontal slice. Contrast: <i>layer</i>.
pattern	A template collaboration.

persistent object	An object that exists after the process or thread that created it has ceased to exist.
postcondition	A constraint that must be true at the completion of an operation.
precondition	A constraint that must be true when an operation is invoked.
primitive type	A pre-defined basic datatype without any substructure, such as an integer or a string.
process	<ol style="list-style-type: none"> 1. A heavyweight unit of concurrency and execution in an operating system. Contrast: <i>thread</i>, which includes heavyweight and lightweight processes. If necessary, an implementation distinction can be made using stereotypes. 2. A software development process—the steps and guidelines by which to develop a system. 3. To execute an algorithm or otherwise handle something dynamically.
projection	A mapping from a set to a subset of it.
property	A named value denoting a characteristic of an element. A property has semantic impact. Certain properties are predefined in the UML; others may be user defined. See: tagged value.
pseudo-state	A vertex in a state machine that has the form of a state, but doesn't behave as a state. Pseudo-states include initial and history vertices.
physical system	<ol style="list-style-type: none"> 1. The subject of a model. 2. A collection of connected physical units, which can include software, hardware and people, that are organized to accomplish a specific purpose. A physical system can be described by one or more models, possibly from different viewpoints. Contrast: system.

published model [MOF]	A model which has been frozen, and becomes available for instantiating repositories and for the support in defining other models. A frozen model's model elements cannot be changed.
qualifier	An association attribute or tuple of attributes whose values partition the set of objects related to an object across an association.
receive [a message]	The handling of a stimulus passed from a sender instance. See: <i>sender, receiver</i> .
receiver [object]	The object handling a stimulus passed from a sender object. Contrast: <i>sender</i> .
reception	A declaration that a classifier is prepared to react to the receipt of a signal.
reference	<ol style="list-style-type: none"> 1. A denotation of a model element. 2. A named slot within a classifier that facilitates navigation to other classifiers. Synonym: <i>pointer</i>.
refinement	A relationship that represents a fuller specification of something that has already been specified at a certain level of detail. For example, a design class is a refinement of an analysis class.
relationship	A semantic connection among model elements. Examples of relationships include associations and generalizations.
repository	A facility for storing object models, interfaces, and implementations.
requirement	A desired feature, property, or behavior of a system.
responsibility	A contract or obligation of a classifier.
reuse	The use of a pre-existing artifact.

role	The named specific behavior of an entity participating in a particular context. A role may be static (e.g., an association end) or dynamic (e.g., a collaboration role).
run time	The period of time during which a computer program executes. Contrast: <i>modeling time</i> .
scenario	A specific sequence of actions that illustrates behaviors. A scenario may be used to illustrate an interaction or the execution of a use case instance. See: <i>interaction</i> .
schema [MOF]	In the context of the MOF, a schema is analogous to a package which is a container of model elements. Schema corresponds to an MOF package. Contrast: <i>metamodel, package</i> .
semantic variation point	A point of variation in the semantics of a metamodel. It provides an intentional degree of freedom for the interpretation of the metamodel semantics.
send [a message]	The passing of a stimulus from a sender instance to a receiver instance. See: <i>sender, receiver</i> .
sender [object]	The object passing a stimulus to a receiver object. Contrast: <i>receiver</i> .
sequence diagram	A diagram that shows object interactions arranged in time sequence. In particular, it shows the objects participating in the interaction and the sequence of messages exchanged. Unlike a collaboration diagram, a sequence diagram includes time sequences but does not include object relationships. A sequence diagram can exist in a generic form (describes all possible scenarios) and in an instance form (describes one actual scenario). Sequence diagrams and collaboration diagrams express similar information, but show it in different ways. See: <i>collaboration diagram</i> .

signal	The specification of an asynchronous stimulus communicated between instances. Signals may have parameters.
signature	The name and parameters of a behavioral feature. A signature may include an optional returned parameter.
single inheritance	A semantic variation of generalization in which a type may have only one supertype. Synonym: <i>multiple inheritance</i> [OMA]. Contrast: <i>multiple inheritance</i> .
single valued [MOF]	A model element with multiplicity defined is single valued when its Multiplicity Type:: upper attribute is set to one. The term single-valued does not pertain to the number of values held by an attribute, parameter, etc., at any point in time, since a single-valued attribute (for instance, with a multiplicity lower bound of zero) may have no value. Contrast: <i>multi-valued</i> .
specification	A declarative description of what something is or does. Contrast: <i>implementation</i> .
state	A condition or situation during the life of an object during which it satisfies some condition, performs some activity, or waits for some event. Contrast: <i>state</i> [OMA].
statechart diagram	A diagram that shows a state machine. See: <i>state machine</i> .
state machine	A behavior that specifies the sequences of states that an object or an interaction goes through during its life in response to events, together with its responses and actions.
static classification	A semantic variation of generalization in which an object may not change classifier. Contrast: <i>dynamic classification</i> .

stereotype	A new type of modeling element that extends the semantics of the metamodel. Stereotypes must be based on certain existing types or classes in the metamodel. Stereotypes may extend the semantics, but not the structure of pre-existing types and classes. Certain stereotypes are predefined in the UML, others may be user defined. Stereotypes are one of three extensibility mechanisms in UML. See: <i>constraint</i> , <i>tagged value</i> .
stimulus	The passing of information from one instance to another, such as raising a signal or invoking an operation. The receipt of a signal is normally considered an event. See: <i>message</i> .
string	A sequence of text characters. The details of string representation depend on implementation, and may include character sets that support international characters and graphics.
structural feature	A static feature of a model element, such as an attribute.
structural model aspect	A model aspect that emphasizes the structure of the objects in a system, including their types, classes, relationships, attributes, and operations.
subactivity state	A state in an activity graph that represents the execution of a non-atomic sequence of steps that has some duration.
subclass	In a generalization relationship, the specialization of another class; the superclass. See: <i>generalization</i> . Contrast: <i>superclass</i> .
submachine state	A state in a state machine which is equivalent to a composite state but its contents is described by another state machine.
substate	A state that is part of a composite state. See: <i>concurrent state</i> , <i>disjoint state</i> .
subpackage	A package that is contained in another package.

subsystem	A grouping of model elements that represents a behavioral unit in a physical system. A subsystem offers interfaces and has operations. In addition, the model elements of a subsystem can be partitioned into specification and realization elements. See <i>package</i> . See: <i>physical system</i> .
subtype	In a generalization relationship, the specialization of another type; the supertype. See: <i>generalization</i> . Contrast: <i>supertype</i> .
superclass	In a generalization relationship, the generalization of another class; the subclass. See: <i>generalization</i> . Contrast: <i>subclass</i> .
supertype	In a generalization relationship, the generalization of another type; the subtype. See: <i>generalization</i> . Contrast: <i>subtype</i> .
supplier	A classifier that provides services that can be invoked by others. Contrast: <i>client</i> .
swimlane	A partition on a activity diagram for organizing the responsibilities for actions. Swimlanes typically correspond to organizational units in a business model. See: <i>partition</i> .
synch state	A vertex in a state machine used for synchronizing the concurrent regions of a state machine.
system	A top-level subsystem in a model. Contrast: <i>physical system</i> .
tagged value	The explicit definition of a property as a name-value pair. In a tagged value, the name is referred as the tag. Certain tags are predefined in the UML; others may be user defined. Tagged values are one of three extensibility mechanisms in UML. See: <i>constraint</i> , <i>stereotype</i> .
template	Synonym: <i>parameterized element</i> .

thread [of control]	A single path of execution through a program, a dynamic model, or some other representation of control flow. Also, a stereotype for the implementation of an active object as lightweight process. See <i>process</i> .
time event	An event that denotes the time elapsed since the current state was entered. See: <i>event</i> .
time expression	An expression that resolves to an absolute or relative value of time.
timing mark	A denotation for the time at which an event or message occurs. Timing marks may be used in constraints.
top level	A stereotype of package denoting the top-most package in a containment hierarchy. The topLevel stereotype defines the outer limit for looking up names, as namespaces “see” outwards. For example, opLevel subsystem represents the top of the subsystem containment hierarchy.
trace	A dependency that indicates a historical or process relationship between two elements that represent the same concept without specific rules for deriving one from the other.
transient object	An object that exists only during the execution of the process or thread that created it.
transition	A relationship between two states indicating that an object in the first state will perform certain specified actions and enter the second state when a specified event occurs and specified conditions are satisfied. On such a change of state, the transition is said to fire.
type	A stereotype of class that is used to specify a domain of instances (objects) together with the operations applicable to the objects. A type may not contain any methods. See: <i>class</i> , <i>instance</i> . Contrast: <i>interface</i> .

type expression	An expression that evaluates to a reference to one or more types.
uninterpreted	A placeholder for a type or types whose implementation is not specified by the UML. Every uninterpreted value has a corresponding string representation. See: <i>any</i> [CORBA].
usage	A dependency in which one element (the client) requires the presence of another element (the supplier) for its correct functioning or implementation.
use case [class]	The specification of a sequence of actions, including variants, that a system (or other entity) can perform, interacting with actors of the system. See: <i>use case instances</i> .
use case diagram	A diagram that shows the relationships among actors and use cases within a system.
use case instance	The performance of a sequence of actions being specified in a use case. An instance of a use case. See: <i>use case class</i> .
use case model	A model that describes a system's functional requirements in terms of use cases.
utility	A stereotype that groups global variables and procedures in the form of a class declaration. The utility attributes and operations become global variables and global procedures, respectively. A utility is not a fundamental modeling construct, but a programming convenience.
value	An element of a type domain.
vertex	A source or a target for a transition in a state machine. A vertex can be either a state or a pseudo-state. See: <i>state</i> , <i>pseudo-state</i> .

view	A projection of a model, which is seen from a given perspective or vantage point and omits entities that are not relevant to this perspective.
view element	A view element is a textual and/or graphical projection of a collection of model elements.
view projection	A projection of model elements onto view elements. A view projection provides a location and a style for each view element.
visibility	An enumeration whose value (public, protected, or private) denotes how the model element to which it refers may be seen outside its enclosing namespace.

The first section of this Appendix gives a rendering of the MOF Model as an XML document encoded using the XML production rules defined in the OMG XMI specification. This XML document is an encoding of the normative MOF Model. The second section of this Appendix gives the normative XMI DTD for MOF meta-model interchange.

The XML document and the DTD were produced automatically using XMI compliant tools.

A.1 The MOF Model in XML

```
<?xml version = '1.0' encoding = 'ISO-8859-1' ?>
<!DOCTYPE XMI SYSTEM 'mof.dtd'>
<XMI xmi.version='1.0'>
  <XMI.header>
    <XMI.metamodel xmi.name='org.omg.mof.Model' xmi.version='1.1' />
  </XMI.header>
  <XMI.content>

  <!-- _____ -->
  <!-- -->
  <!-- Contents of Package: Model -->
  <!-- -->
  <!-- _____ -->

  <Model.Package xmi.id='a1'>
    <Model.ModelElement.name>Model</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
```

```
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.Namespace.contents>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: NameType -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a97'>
  <Model.ModelElement.name>NameType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcAlias xmi.tcName='NameType'
        xmi.tcId='IDL:org.omg.mof/Model/NameType:1.0'>
        <XMI.CorbaTypeCode>
          <XMI.CorbaTcString xmi.tcLength='0' />
        </XMI.CorbaTypeCode>
      </XMI.CorbaTcAlias>
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: AnnotationType -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a109'>
  <Model.ModelElement.name>AnnotationType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
```

```

<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.DataType.typeCode>
  <XMI.CorbaTypeCode>
    <XMI.CorbaTcAlias xmi.tcName='AnnotationType'
      xmi.tcId='IDL:org.omg.mof/Model/AnnotationType:1.0'>
      <XMI.CorbaTypeCode>
        <XMI.CorbaTcString xmi.tcLength='0' />
      </XMI.CorbaTypeCode>
    </XMI.CorbaTcAlias>
  </XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: ModelElement -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a29'>
  <Model.ModelElement.name>ModelElement</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='true' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c1' />
    <Model.Constraint xmi.idref='c2' />
    <Model.Constraint xmi.idref='c3' />
    <Model.Constraint xmi.idref='c4' />
  </Model.ModelElement.constraints>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c1'>
      <Model.ModelElement.name>MustBeContainedUnlessPackage</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">

```

```

    inv:
    not self.oclIsTypeOf(Package) implies
    self.container -&gt; size = 1
</XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='deferred' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c2'>
  <Model.ModelElement.name>FrozenAttributesCannotBeChanged</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
  <XMI.any xmi.type="string" xmi.name="">
    inv:
    self.isFrozen() implies
    let myTypes = self.oclType() -&gt; allSupertypes() -&gt;
    includes(self.oclType()) in
    let myAttrs : Set(Attribute) =
    self.RefBaseObject::refMetaObject() -&gt;
    asOclType(Class) -&gt;
    findElementsByTypeExtended(Attribute) in
    myAttrs -&gt; forAll(a |
    self.RefObject::refValue@pre(a) =
    self.RefObject::refValue(a))
  </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c3'>
  <Model.ModelElement.name>FrozenElementsCannotBeDeleted</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
  <XMI.any xmi.type="string" xmi.name="">
    post:

```

```

    (self.isFrozen@pre() and
    self.container@pre -&gt; notEmpty and
    self.container.isFrozen@pre()) implies
    (self.container.Object::non_existent() or
    not self.Object::non_existent())
</XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c4'>
  <Model.ModelElement.name>FrozenDependenciesCannotBeChanged</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      post:
        self.isFrozen() implies
        let myClasses = self.oclType() -&gt; allSupertypes() -&gt;
        includes(self.oclType()) in
        let myRefs = Set(Reference) =
        self.RefBaseObject::refMetaObject() -&gt;
        asOclType(Class) -&gt;
        findElementsByTypeExtended(Reference) in
        let myDepRefs = myRefs -&gt;
        select(r |
        Set{"contents", "constraints", "supertypes",
        "type", "referencedEnd", "exceptions",
        "importedNamespace", "elements"} -&gt;
        includes(r.name)) in
        myDepRefs -&gt;
        forAll(r |
        self.RefObject::refValue@pre(r) =
        self.RefObject::refValue(r))
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>

```

```
<Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>name</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>qualifiedName</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>true</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='false' />
  <Model.Attribute.isDerived xmi.value='true' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>annotation</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
```

```

<Model.Feature.scope xmi.value='instance_level' />
<Model.StructuralFeature.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.Attribute.isDerived xmi.value='false' />
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a109' /> <!-- Model.AnnotationType -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Reference>
  <Model.ModelElement.name>requiredElements</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='false' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a131' /> <!-- Model.DependsOn.provider -->
  </Model.Reference.referencedEnd>
</Model.Reference>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: DependencyKind -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a116'>
  <Model.ModelElement.name>DependencyKind</Model.ModelElement.name>

```

```

<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.DataType.typeCode>
  <XMI.CorbaTypeCode>
    <XMI.CorbaTcAlias xmi.tcName='DependencyKind'
      xmi.tcId='IDL:org.omg.mof/Model/ModelElementClass/DependencyKind:1.0'>
      <XMI.CorbaTypeCode>
        <XMI.CorbaTcString xmi.tcLength='0' />
      </XMI.CorbaTypeCode>
    </XMI.CorbaTcAlias>
  </XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<Model.Constant>
  <Model.ModelElement.name>ContainerDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">container</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>

<Model.Constant>
  <Model.ModelElement.name>ContentsDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">contents</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>

<Model.Constant>
  <Model.ModelElement.name>SignatureDep</Model.ModelElement.name>

```



```

    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constant.value>
      <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">signature</XMI.any>
    </Model.Constant.value>
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
    </Model.TypedElement.type>
  </Model.Constant>
  <Model.Constant>
    <Model.ModelElement.name>ConstraintDep</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constant.value>
      <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">constraint</XMI.any>
    </Model.Constant.value>
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
    </Model.TypedElement.type>
  </Model.Constant>
  <Model.Constant>
    <Model.ModelElement.name>ConstrainedElementsDep</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constant.value>
      <XMI.any xmi.type="string" xmi.name="Model.ModelElement.DependencyKind">constrained
elements</XMI.any>
    </Model.Constant.value>
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
    </Model.TypedElement.type>
  </Model.Constant>
  <Model.Constant>
    <Model.ModelElement.name>SpecializationDep</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constant.value>
      <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">specialization</XMI.any>
    </Model.Constant.value>
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
    </Model.TypedElement.type>
  </Model.Constant>

```

```

<Model.Constant>
  <Model.ModelElement.name>ImportDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">import</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>
<Model.Constant>
  <Model.ModelElement.name>TypeDefinitionDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string" xmi.name="Model.ModelElement.DependencyKind">type
definition</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>
<Model.Constant>
  <Model.ModelElement.name>ReferencedEndsDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string" xmi.name="Model.ModelElement.DependencyKind">referenced
ends</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>
<Model.Constant>
  <Model.ModelElement.name>TaggedElementsDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string" xmi.name="Model.ModelElement.DependencyKind">tagged
elements</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->

```

```

    </Model.TypedElement.type>
</Model.Constant>
<Model.Constant>
  <Model.ModelElement.name>IndirectDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string"
xmi.name="Model.ModelElement.DependencyKind">indirect</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>
<Model.Constant>
  <Model.ModelElement.name>AllDep</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="string" xmi.name="Model.ModelElement.DependencyKind">all</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Constant>
<Model.Operation>
  <Model.ModelElement.name>findRequiredElements</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>

```

```

    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>kinds</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>recursive</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>isRequiredBecause</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>

```

```

<Model.ModelElement.name>*return</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.Parameter.direction xmi.value='return_dir' />
<Model.Parameter.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
</Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>otherElement</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>reason</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='out_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a116' /> <!-- ModelElement.DependencyKind -->
  </Model.TypedElement.type>
</Model.Parameter>

```

```
</Model.Namespace.contents>
</Model.Operation>
<Model.Reference>
  <Model.ModelElement.name>container</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a120' /> <!-- Model.Contains.container -->
  </Model.Reference.referencedEnd>
</Model.Reference>
<Model.Reference>
  <Model.ModelElement.name>constraints</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a58' /> <!-- Model.Constraint -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a126' /> <!-- Model.Constrains.constraint -->
  </Model.Reference.referencedEnd>
</Model.Reference>
```

```

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: VerifyResultKind -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a117'>
  <Model.ModelElement.name>VerifyResultKind</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcEnum xmi.tcName='VerifyResultKind'
        xmi.tcId='IDL:org.omg.mof/Model/ModelElementClass/VerifyResultKind:1.0'>
        <XMI.CorbaTcEnumLabel xmi.tcName='valid' />
        <XMI.CorbaTcEnumLabel xmi.tcName='invalid' />
        <XMI.CorbaTcEnumLabel xmi.tcName='published' />
      </XMI.CorbaTcEnum>
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: DepthKind -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a111'>
  <Model.ModelElement.name>DepthKind</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcEnum xmi.tcName='DepthKind'

```

```

        xmi.tcId='IDL:org.omg.mof/Model/ModelElementClass/DepthKind:1.0'>
    <XMI.CorbaTcEnumLabel xmi.tcName='shallow' />
    <XMI.CorbaTcEnumLabel xmi.tcName='deep' />
</XMI.CorbaTcEnum>
</XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: ViolationType -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a103'>
  <Model.ModelElement.name>ViolationType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcStruct xmi.tcName='ViolationType'
        xmi.tcId='IDL:org.omg.mof/Model/ModelElementClass/ViolationType:1.0'>
        <XMI.CorbaTcField xmi.tcName='error_kind'>
          <XMI.CorbaTypeCode>
            <XMI.CorbaTcString xmi.tcLength='0' />
          </XMI.CorbaTypeCode>
        </XMI.CorbaTcField>
        <XMI.CorbaTcField xmi.tcName='element_in_error'>
          <XMI.CorbaTypeCode>
            <XMI.CorbaTcObjRef xmi.tcName=''
xmi.tcId='IDL:org.omg.mof/Reflective/RefObject:1.0' />
          </XMI.CorbaTypeCode>
        </XMI.CorbaTcField>
        <XMI.CorbaTcField xmi.tcName='values_in_error'>
          <XMI.CorbaTypeCode>
            <XMI.CorbaTcObjRef xmi.tcName=''
xmi.tcId='IDL:org.omg.mof/Reflective/NamedObjectList:1.0' />
          </XMI.CorbaTypeCode>
        </XMI.CorbaTcField>
      </XMI.CorbaTypeCode>
    </Model.DataType.typeCode>
  </Model.DataType>

```



```

    </XMI.CorbaTcField>
    <XMI.CorbaTcField xmi.tcName='error_description'>
      <XMI.CorbaTypeCode>
        <XMI.CorbaTcString xmi.tcLength='0' />
      </XMI.CorbaTypeCode>
    </XMI.CorbaTcField>
  </XMI.CorbaTcStruct>
</XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>
<Model.Operation>
  <Model.ModelElement.name>verify</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a117' /> <!-- ModelElement.VerifyResultKind -->
      </Model.TypedElement.type>
    </Model.Parameter>
    <Model.Parameter>
      <Model.ModelElement.name>depth</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='in_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>

```

```
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a111' /> <!-- ModelElement.DepthKind -->
</Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>violations</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='out_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a103' /><!-- ModelElement.ViolationType -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>isFrozen</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
      </Model.TypedElement.type>
    </Model.Parameter>
```

```

    </Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>isVisible</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
      </Model.TypedElement.type>
    </Model.Parameter>
    <Model.Parameter>
      <Model.ModelElement.name>otherElement</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='in_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
      </Model.TypedElement.type>
    </Model.Parameter>
  </Model.Namespace.contents>
</Model.Operation>
</Model.Namespace.contents>
</Model.Class>

```

```
<!-- _____ -->
<!-- -->
<!-- Contents of DataType: VisibilityKind -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a110'>
  <Model.ModelElement.name>VisibilityKind</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcEnum xmi.tcName='VisibilityKind'
        xmi.tcId='IDL:org.omg.mof/Model/VisibilityKind:1.0'>
        <XMI.CorbaTcEnumLabel xmi.tcName='public_vis' />
        <XMI.CorbaTcEnumLabel xmi.tcName='protected_vis' />
        <XMI.CorbaTcEnumLabel xmi.tcName='private_vis' />
      </XMI.CorbaTcEnum>
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Namespace -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a6'>
  <Model.ModelElement.name>Namespace</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='true' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
```

```

    <Model.Constraint xmi.idref='c5' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c5'>
      <Model.ModelElement.name>ContentNamesMustNotCollide</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv: self.contents.forAll(
            e1, e2 | e1.name = e2.name implies r1 = r2)
        </XMI.any>
      </Model.Constraint.expression>
      <Model.Constraint.language>OCL</Model.Constraint.language>
      <Model.Constraint.evaluationPolicy xmi.value='immediate' />
      <Model.Constraint.constrainedElements>
        <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
      </Model.Constraint.constrainedElements>
    </Model.Constraint>
    <Model.Exception xmi.id='NameNotFound'>
      <Model.ModelElement.name>NameNotFound</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Feature.visibility xmi.value='public_vis' />
      <Model.Feature.scope xmi.value='instance_level' />
      <Model.Namespace.contents>
        <Model.Parameter>
          <Model.ModelElement.name>name</Model.ModelElement.name>
          <Model.ModelElement.annotation></Model.ModelElement.annotation>
          <Model.Parameter.direction xmi.value='out_dir' />
          <Model.Parameter.multiplicity>
            <XMI.field>1</XMI.field>
            <XMI.field>1</XMI.field>
            <XMI.field>>false</XMI.field>
            <XMI.field>>false</XMI.field>
          </Model.Parameter.multiplicity>
          <Model.TypedElement.type>
            <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
          </Model.TypedElement.type>
        </Model.Parameter>
      </Model.Namespace.contents>
    </Model.Exception>
  </Model.Namespace.contents>

```

```
</Model.Namespace.contents>
</Model.Exception>
<Model.Exception xmi.id='NameNotResolved'>
  <Model.ModelElement.name>NameNotResolved</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>explanation</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='out_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a100' /> <!-- Model.string -->
      </Model.TypedElement.type>
    </Model.Parameter>
    <Model.Parameter>
      <Model.ModelElement.name>restOfName</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='out_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>true</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
      </Model.TypedElement.type>
    </Model.Parameter>
  </Model.Namespace.contents>
</Model.Exception>
<Model.Reference>
  <Model.ModelElement.name>contents</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
```

```

<Model.Feature.visibility xmi.value='public_vis' />
<Model.Feature.scope xmi.value='instance_level' />
<Model.StructuralFeature.multiplicity>
  <XMI.field>0</XMI.field>
  <XMI.field>-1</XMI.field>
  <XMI.field>>true</XMI.field>
  <XMI.field>>true</XMI.field>
</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.TypedElement.type>
  <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.TypedElement.type>
<Model.Reference.referencedEnd>
  <Model.AssociationEnd xmi.idref='a121' /> <!-- Model.Contains.containedElement -->
</Model.Reference.referencedEnd>
</Model.Reference>
<Model.Operation>
  <Model.ModelElement.name>lookupElement</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Operation.exceptions>
    <Model.Exception xmi.idref='NameNotFound' />
  </Model.Operation.exceptions>
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
      </Model.TypedElement.type>
    </Model.Parameter>
  <Model.Parameter>

```

```

<Model.ModelElement.name>name</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.Parameter.direction xmi.value='in_dir' />
<Model.Parameter.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
</Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>resolveQualifiedNames</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Operation.exceptions>
    <Model.Exception xmi.idref='NameNotResolved' />
  </Model.Operation.exceptions>
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
      </Model.TypedElement.type>
    </Model.Parameter>
  </Model.Namespace.contents>
  <Model.Parameter>
    <Model.ModelElement.name>qualifiedName</Model.ModelElement.name>

```



```

<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.Parameter.direction xmi.value='in_dir' />
<Model.Parameter.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>-1</XMI.field>
  <XMI.field>>true</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
</Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
<Model.ModelElement.name>findElementsByType</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.Feature.visibility xmi.value='public_vis' />
<Model.Feature.scope xmi.value='instance_level' />
<Model.Operation.isQuery xmi.value='true' />
<Model.Namespace.contents>
<Model.Parameter>
  <Model.ModelElement.name>*return</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='return_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>true</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>ofType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>

```

```
<XMI.field>1</XMI.field>
<XMI.field>>false</XMI.field>
<XMI.field>>false</XMI.field>
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.Class xmi.idref='a27' /> <!-- Model.Class -->
</Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>includeSubtypes</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>nameIsValid</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
```

```

    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
    </Model.TypedElement.type>
  </Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>proposedName</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: GeneralizableElement -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a71'>
  <Model.ModelElement.name>GeneralizableElement</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='true' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c6' />
    <Model.Constraint xmi.idref='c7' />
    <Model.Constraint xmi.idref='c8' />

```

```
<Model.Constraint xmi.idref='c9' />
<Model.Constraint xmi.idref='c10' />
<Model.Constraint xmi.idref='c11' />
<Model.Constraint xmi.idref='c12' />
</Model.ModelElement.constraints>
<Model.GeneralizableElement.supertypes>
  <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
</Model.GeneralizableElement.supertypes>
<Model.Namespace.contents>
  <Model.Constraint xmi.id='c6'>
    <Model.ModelElement.name>SupertypeMustNotBeSelf</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.allSupertypes() -&gt; forAll(s | s &lt;&gt; self)
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c7'>
    <Model.ModelElement.name>SupertypeKindMustBeSame</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.supertypes -&gt; forAll(s | s.ocType() = self.ocType())
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c8'>
    <Model.ModelElement.name>ContentsMustNotCollideWithSupertypes</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
```

```

<XMI.any xmi.type="string" xmi.name="">
  inv:
  let superContents = self.allSupertypes() -&gt;
  collect(s | s.contents) in
  self.contents -&gt;
  forAll(m1 |
  superContents -&gt;
  forAll(m2 |
  m1.name = m2.name implies m1 = m2))
</XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c9'>
  <Model.ModelElement.name>DiamondRuleMustBeObeyed</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
  <XMI.any xmi.type="string" xmi.name="">
    inv:
    let superNamespaces =
    self.supertypes -&gt; collect(s | s.extendedNamespace) in
    superNamespaces -&gt; asSet -&gt; isUnique(s | s.name)
  </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c10'>
  <Model.ModelElement.name>NoSupertypesAllowedForRoot</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
  <XMI.any xmi.type="string" xmi.name="">
    inv: self.isRoot implies self.supertypes -&gt; isEmpty
  </XMI.any>

```

```

</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c11'>
  <Model.ModelElement.name>SupertypesMustBeVisible</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.supertypes -&gt; forAll(s | self.isVisible(s))
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c12'>
  <Model.ModelElement.name>NoSubtypesAllowedForLeaf</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.supertypes -&gt; forAll(s | not s.isLeaf)
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>isRoot</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>

```

```

    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>isLeaf</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>isAbstract</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />

```

```

<Model.TypedElement.type>
  <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>visibility</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a110' /> <!-- Model.VisibilityKind -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Reference xmi.id='a140'>
  <Model.ModelElement.name>supertypes</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>true</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a138' /> <!-- Model.Generalizes.supertype -->
  </Model.Reference.referencedEnd>
</Model.Reference>
<Model.Operation>

```



```

<Model.ModelElement.name>allSupertypes</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.Feature.visibility xmi.value='public_vis' />
<Model.Feature.scope xmi.value='instance_level' />
<Model.Operation.isQuery xmi.value='true' />
<Model.Namespace.contents>
  <Model.Parameter>
    <Model.ModelElement.name>*return</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Parameter.direction xmi.value='return_dir' />
    <Model.Parameter.multiplicity>
      <XMI.field>0</XMI.field>
      <XMI.field>-1</XMI.field>
      <XMI.field>>true</XMI.field>
      <XMI.field>>true</XMI.field>
    </Model.Parameter.multiplicity>
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
    </Model.TypedElement.type>
  </Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>lookupElementExtended</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Operation.exceptions>
    <Model.Exception xmi.idref='NameNotFound' />
  </Model.Operation.exceptions>
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
    </Model.Parameter>
  </Model.Namespace.contents>
</Model.Operation>

```

```
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>name</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a97' /> <!-- Model.NameType -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>findElementsByTypeExtended</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>true</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
      </Model.TypedElement.type>
```

```

</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>ofType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a27' /> <!-- Model.Class -->
  </Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>includeSubtypes</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of Class: TypedElement -->
<!-- _____ -->
<!-- _____ -->

<Model.Class xmi.id='a96'>

```

```
<Model.ModelElement.name>TypedElement</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='true' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.Class.isSingleton xmi.value='false' />
<Model.ModelElement.constraints>
  <Model.Constraint xmi.idref='c13' />
  <Model.Constraint xmi.idref='c14' />
</Model.ModelElement.constraints>
<Model.GeneralizableElement.supertypes>
  <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.GeneralizableElement.supertypes>
<Model.Namespace.contents>
  <Model.Constraint xmi.id='c13'>
    <Model.ModelElement.name>AssociationsCannotBeTypes</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: not self.type.ocIsKindOf(Association)
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c14'>
    <Model.ModelElement.name>TypeMustBeVisible</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.isVisible(self.type)
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='deferred' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
</Model.Namespace.contents>
</Model.Namespace>
```

```

    </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Reference>
  <Model.ModelElement.name>type</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a5' /> <!-- Model.Classifier -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a155' /> <!-- Model.IsOfType.type -->
  </Model.Reference.referencedEnd>
</Model.Reference>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Classifier -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a5'>
  <Model.ModelElement.name>Classifier</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='true' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
  </Model.GeneralizableElement.supertypes>

```

```

</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Class -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a27'>
  <Model.ModelElement.name>Class</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c15' />
    <Model.Constraint xmi.idref='c16' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a5' /> <!-- Model.Classifier -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c15'>
      <Model.ModelElement.name>ClassContainmentRules</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv:
            Set{Class, DataType, Attribute, Reference, Operation,
              Exception, Constraint, Tag} -&gt;
            includesAll(self.contentTypes())
        </XMI.any>
      </Model.Constraint.expression>
      <Model.Constraint.language>OCL</Model.Constraint.language>
      <Model.Constraint.evaluationPolicy xmi.value='immediate' />
      <Model.Constraint.constrainedElements>
        <Model.Class xmi.idref='a27' /> <!-- Model.Class -->
      </Model.Constraint.constrainedElements>
    </Model.Constraint>
  </Model.Namespace.contents>
</Model.Class>

```

```

<Model.Constraint xmi.id='c16'>
  <Model.ModelElement.name>AbstractClassesCannotBeSingleton</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.isAbstract implies not self.isSingleton
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a27' /> <!-- Model.Class -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>isSingleton</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Attribute>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of DataType: TypeDescriptor -->
<!-- _____ -->
<!-- _____ -->

<Model.DataType xmi.id='a102'>

```

```

<Model.ModelElement.name>TypeDescriptor</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.DataType.typeCode>
  <XMI.CorbaTypeCode>
    <XMI.CorbaTcAlias xmi.tcName='TypeDescriptor'
      xmi.tcId='IDL:org.omg.mof/Model/TypeDescriptor:1.0'>
      <XMI.CorbaTypeCode>
        <XMI.CorbaTcTypeCode />
      </XMI.CorbaTypeCode>
    </XMI.CorbaTcAlias>
  </XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of Class: DataType _____ -->
<!-- _____ -->
<!-- _____ -->

<Model.Class xmi.id='a88'>
  <Model.ModelElement.name>DataType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c17' />
    <Model.Constraint xmi.idref='c18' />
    <Model.Constraint xmi.idref='c19' />
    <Model.Constraint xmi.idref='c20' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a5' /> <!-- Model.Classifier -->
  </Model.GeneralizableElement.supertypes>

```



```

<Model.Namespace.contents>
  <Model.Constraint xmi.id='c17'>
    <Model.ModelElement.name>DataTypeContainmentRules</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv:
          Set{TypeAlias, Constraint, Tag} -&gt;
            includesAll(self.contentTypes())
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a27' /> <!-- Model.Class -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c18'>
    <Model.ModelElement.name>ThisTypecodeNotSupported</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv:
          self.typeCode.allTypeKinds() -&gt;
            excludes(Set{#tk_void, #tk_Principal, #tk_null, #tk_except,
              #tk_value, #tk_value_box, #tk_native,
              #tk_abstract_interface})
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='deferred' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a88' /> <!-- Model.DataType -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c19'>
    <Model.ModelElement.name>DataTypesHaveNoSupertypes</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.supertypes -&gt; isEmpty
      </XMI.any>
    </Model.Constraint.expression>
  </Model.Constraint>

```

```
</XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a88' /> <!-- Model.DataType -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c20'>
  <Model.ModelElement.name>DataTypesCannotBeAbstract</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: not self.isAbstract
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a88' /> <!-- Model.DataType -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>typeCode</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a102' /> <!-- Model.TypeDescriptor -->
  </Model.TypedElement.type>
</Model.Attribute>
</Model.Namespace.contents>
</Model.Class>
```

```

<!-- _____ -->
<!-- -->
<!-- Contents of Class: TypeAlias -->
<!-- -->
<!-- _____ -->

```

```

<Model.Class xmi.id='a104'>
  <Model.ModelElement.name>TypeAlias</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
  </Model.GeneralizableElement.supertypes>
</Model.Class>

```

```

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: ScopeKind -->
<!-- -->
<!-- _____ -->

```

```

<Model.DataType xmi.id='a113'>
  <Model.ModelElement.name>ScopeKind</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcEnum xmi.tcName='ScopeKind'
        xmi.tcId='IDL:org.omg.Mof/Model/ScopeKind:1.0'>
        <XMI.CorbaTcEnumLabel xmi.tcName='instance_level' />
        <XMI.CorbaTcEnumLabel xmi.tcName='classifier_level' />
      </XMI.CorbaTcEnum>
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>

```

```

    </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Feature -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a2'>
  <Model.ModelElement.name>Feature</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='true' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Attribute>
      <Model.ModelElement.name>scope</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Feature.visibility xmi.value='public_vis' />
      <Model.Feature.scope xmi.value='instance_level' />
      <Model.StructuralFeature.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.StructuralFeature.multiplicity>
      <Model.StructuralFeature.isChangeable xmi.value='true' />
      <Model.Attribute.isDerived xmi.value='false' />
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a113' /> <!-- Model.ScopeKind -->
      </Model.TypedElement.type>
    </Model.Attribute>
    <Model.Attribute>
      <Model.ModelElement.name>visibility</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>

```

```

    <Model.Feature.visibility xmi.value='public_vis' />
    <Model.Feature.scope xmi.value='instance_level' />
    <Model.StructuralFeature.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>false</XMI.field>
    </Model.StructuralFeature.multiplicity>
    <Model.StructuralFeature.isChangeable xmi.value='true' />
    <Model.Attribute.isDerived xmi.value='false' />
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a110' /> <!-- Model.VisibilityKind -->
    </Model.TypedElement.type>
  </Model.Attribute>
</Model.Namespace.contents>
</Model.Class>

<Model.Constant>
  <Model.ModelElement.name>Unbounded</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constant.value>
    <XMI.any xmi.type="uLong" xmi.name="">-1</XMI.any>
  </Model.Constant.value>
  <Model.TypedElement.type>
    <Model.Class xmi.idref='uLong' /> <!-- unsigned long -->
  </Model.TypedElement.type>
</Model.Constant>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: MultiplicityType -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a99'>
  <Model.ModelElement.name>MultiplicityType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />

```

```
<Model.DataType.typeCode>
  <XMI.CorbaTypeCode>
    <XMI.CorbaTcStruct xmi.tcName='MultiplicityType'
      xmi.tcId='IDL:org.omg.mof/Model/MultiplicityType:1.0'>
      <XMI.CorbaTcField xmi.tcName='lower'>
        <XMI.CorbaTypeCode>
          <XMI.CorbaTcLong/>
        </XMI.CorbaTypeCode>
      </XMI.CorbaTcField>
      <XMI.CorbaTcField xmi.tcName='upper'>
        <XMI.CorbaTypeCode>
          <XMI.CorbaTcLong/>
        </XMI.CorbaTypeCode>
      </XMI.CorbaTcField>
      <XMI.CorbaTcField xmi.tcName='is_ordered'>
        <XMI.CorbaTypeCode>
          <XMI.CorbaTcBoolean/>
        </XMI.CorbaTypeCode>
      </XMI.CorbaTcField>
      <XMI.CorbaTcField xmi.tcName='is_unique'>
        <XMI.CorbaTypeCode>
          <XMI.CorbaTcBoolean/>
        </XMI.CorbaTypeCode>
      </XMI.CorbaTcField>
    </XMI.CorbaTcStruct>
  </XMI.CorbaTypeCode>
</Model.DataType.typeCode>
<Model.ModelElement.constraints>
  <Model.Constraint xmi.idref='c55' />
  <Model.Constraint xmi.idref='c56' />
  <Model.Constraint xmi.idref='c57' />
  <Model.Constraint xmi.idref='c58' />
</Model.ModelElement.constraints>
<Model.Namespace.contents>
  <Model.Constraint xmi.id='c55'>
    <Model.ModelElement.name>LowerCannotBeNegativeOrUnbounded</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.lower &gt;= 0 and self.lower &lt;&gt; Unbounded
      </XMI.any>
    </Model.Constraint.expression>
  </Model.Constraint>
</Model.Namespace.contents>
```

```

</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a99' /> <!-- Model.MultiplicityType -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c56'>
  <Model.ModelElement.name>LowerCannotExceedUpper</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.lower &lt;= self.upper or self.upper = Unbounded
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a99' /> <!-- Model.MultiplicityType -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c57'>
  <Model.ModelElement.name>UpperMustBePositive</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.upper &gt;= 1 or self.upper = Unbounded
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a99' /> <!-- Model.MultiplicityType -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c58'>
  <Model.ModelElement.name>MustBeUnorderedNonunique</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:

```

```

    self.upper = 1 implies
    (not self.isOrdered and not self.isUnique)
  </XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a99' /> <!-- Model.MultiplicityType -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
</Model.Namespace.contents>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: StructuralFeature -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a20'>
  <Model.ModelElement.name>StructuralFeature</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='true' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a2' /> <!-- Model.Feature -->
    <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Attribute>
      <Model.ModelElement.name>multiplicity</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Feature.visibility xmi.value='public_vis' />
      <Model.Feature.scope xmi.value='instance_level' />
      <Model.StructuralFeature.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>

```



```

    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a99' /> <!-- Model.MultiplicityType -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>isChangeable</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Attribute>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of Class: Attribute -->
<!-- _____ -->
<!-- _____ -->

<Model.Class xmi.id='a23'>
  <Model.ModelElement.name>Attribute</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />

```

```

<Model.Class.isSingleton xmi.value='false' />
<Model.GeneralizableElement.supertypes>
  <Model.Class xmi.idref='a20' /> <!-- Model.StructuralFeature -->
</Model.GeneralizableElement.supertypes>
<Model.Namespace.contents>
  <Model.Attribute>
    <Model.ModelElement.name>isDerived</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Feature.visibility xmi.value='public_vis' />
    <Model.Feature.scope xmi.value='instance_level' />
    <Model.StructuralFeature.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>false</XMI.field>
    </Model.StructuralFeature.multiplicity>
    <Model.StructuralFeature.isChangeable xmi.value='true' />
    <Model.Attribute.isDerived xmi.value='false' />
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
    </Model.TypedElement.type>
  </Model.Attribute>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Reference -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a86'>
  <Model.ModelElement.name>Reference</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c21' />

```

```

<Model.Constraint xmi.idref='c22' />
<Model.Constraint xmi.idref='c23' />
<Model.Constraint xmi.idref='c24' />
<Model.Constraint xmi.idref='c25' />
<Model.Constraint xmi.idref='c26' />
<Model.Constraint xmi.idref='c27' />
</Model.ModelElement.constraints>
<Model.GeneralizableElement.supertypes>
  <Model.Class xmi.idref='a20' /> <!-- Model.StructuralFeature -->
</Model.GeneralizableElement.supertypes>
<Model.Namespace.contents>
  <Model.Constraint xmi.id='c21'>
    <Model.ModelElement.name>ReferenceMultiplicityMustMatchEnd</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.multiplicity = self.referencedEnd.multiplicity
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='deferred' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c22'>
    <Model.ModelElement.name>ReferenceMustBeInstanceScoped</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: self.scope = #instance_level
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c23'>

```

```
<Model.ModelElement.name>ChangeableReferenceMustHaveChangeableEnd</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.isChangeable = self.referencedEnd.isChangeable
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c24'>
  <Model.ModelElement.name>ReferenceTypeMustMatchEndType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.type = self.referencedEnd.type
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c25'>
  <Model.ModelElement.name>ReferencedEndMustBeNavigable</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.referencedEnd.isNavigable
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
  </Model.Constraint.constrainedElements>
```

```

</Model.Constraint>
<Model.Constraint xmi.id='c26'>
  <Model.ModelElement.name>ContainerMustMatchExposedType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
        self.container.allSupertypes() -&gt; including(self) -&gt;
          includes(self.referencedEnd.otherEnd.type)
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c27'>
  <Model.ModelElement.name>ReferencedEndMustBeVisible</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.isVisible(self.referencedEnd)
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Reference>
  <Model.ModelElement.name>exposedEnd</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>

```

```

</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.TypedElement.type>
  <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
</Model.TypedElement.type>
<Model.Reference.referencedEnd>
  <Model.AssociationEnd xmi.idref='a150' /> <!-- Model.Exposes.exposedEnd -->
</Model.Reference.referencedEnd>
</Model.Reference>
<Model.Reference>
  <Model.ModelElement.name>referencedEnd</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a146' /> <!-- Model.RefersTo.referencedEnd -->
  </Model.Reference.referencedEnd>
</Model.Reference>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of Class: BehavioralFeature -->
<!-- _____ -->
<!-- _____ -->

<Model.Class xmi.id='a67'>
  <Model.ModelElement.name>BehavioralFeature</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />

```

```

<Model.GeneralizableElement.isAbstract xmi.value='true' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.Class.isSingleton xmi.value='false' />
<Model.GeneralizableElement.supertypes>
  <Model.Class xmi.idref='a2' /> <!-- Model.Feature -->
  <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
</Model.GeneralizableElement.supertypes>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Operation -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a25'>
  <Model.ModelElement.name>Operation</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c28' />
    <Model.Constraint xmi.idref='c29' />
    <Model.Constraint xmi.idref='c30' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a67' /> <!-- Model.BehavioralFeature -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c28'>
      <Model.ModelElement.name>OperationContainmentRules</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv:
          Set{Parameter, Constraint, Tag} -&gt;
          includesAll(self.contentTypes())
        </XMI.any>
      </Model.Constraint.expression>
    </Model.Constraint>
  </Model.Namespace.contents>

```

```

    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a25' /> <!-- Model.Operation -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c29'>
  <Model.ModelElement.name>OperationsHaveAtMostOneReturn</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
        self.contents -&gt;
          select(c | c.oclIsTypeOf(Parameter)) -&gt;
            select(p : Parameter | p.direction = #return_dir) -&gt;
              size &lt; 2
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a25' /> <!-- Model.Operation -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c30'>
  <Model.ModelElement.name>OperationExceptionsMustBeVisible</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.exceptions -&gt; forAll(e | self.isVisible(e))
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a25' /> <!-- Model.Operation -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>

```



```

<Model.ModelElement.name>isQuery</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.Feature.visibility xmi.value='public_vis' />
<Model.Feature.scope xmi.value='instance_level' />
<Model.StructuralFeature.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.Attribute.isDerived xmi.value='false' />
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Reference xmi.id='a136'>
  <Model.ModelElement.name>exceptions</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>true</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a87' /> <!-- Model.Exception -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a134' /> <!-- Model.CanRaise.except -->
  </Model.Reference.referencedEnd>
</Model.Reference>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Exception -->

```

```

<!--                                     -->
<!-- _____ -->

<Model.Class xmi.id='a87'>
  <Model.ModelElement.name>Exception</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c31' />
    <Model.Constraint xmi.idref='c32' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a67' /> <!-- Model.BehavioralFeature -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c31'>
      <Model.ModelElement.name>ExceptionContainmentRules</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv: Set{Parameter, Tag} -&gt; includesAll(self.contentTypes())
        </XMI.any>
      </Model.Constraint.expression>
      <Model.Constraint.language>OCL</Model.Constraint.language>
      <Model.Constraint.evaluationPolicy xmi.value='immediate' />
      <Model.Constraint.constrainedElements>
        <Model.Class xmi.idref='a87' /> <!-- Model.Exception -->
      </Model.Constraint.constrainedElements>
    </Model.Constraint>
    <Model.Constraint xmi.id='c32'>
      <Model.ModelElement.name>ExceptionsHaveOnlyOutParameters</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv:
            self.contents -&gt;
              select(c | c.oclIsTypeOf(Parameter)) -&gt;

```

```

    forAll(p : Parameter | p.direction = #out_dir)
  </XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a87' /> <!-- Model.Exception -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Association -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a62'>
  <Model.ModelElement.name>Association</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c33' />
    <Model.Constraint xmi.idref='c34' />
    <Model.Constraint xmi.idref='c36' />
    <Model.Constraint xmi.idref='c37' />
    <Model.Constraint xmi.idref='c38' />
    <Model.Constraint xmi.idref='c39' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a5' /> <!-- Model.Classifier -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c33'>
      <Model.ModelElement.name>AssociationContainmentRules</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>

```

```
<Model.Constraint.expression>
  <XMI.any xmi.type="string" xmi.name="">
    inv:
      Set{AssociationEnd, Constraint, Tag} -&gt;
        includesAll(self.contentTypes())
  </XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='immediate' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a62' /> <!-- Model.Association -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c34'>
  <Model.ModelElement.name>AssociationsHaveNoSupertypes</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.supertypes -&gt; isEmpty
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a62' /> <!-- Model.Association -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c36'>
  <Model.ModelElement.name>AssociationMustBeRootAndLeaf</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.isRoot and self.isLeaf
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a62' /> <!-- Model.Association -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
```

```

<Model.Constraint xmi.id='c37'>
  <Model.ModelElement.name>AssociationsCannotBeAbstract</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: not self.isAbstract
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a62' /> <!-- Model.Association -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c38'>
  <Model.ModelElement.name>AssociationsMustBePublic</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.visibility = #public_vis
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a62' /> <!-- Model.Association -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c39'>
  <Model.ModelElement.name>AssociationsMustBeBinary</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.contents -&gt;
        select(c | c.oclIsTypeOf(AssociationEnd)) -&gt; size = 2
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a62' /> <!-- Model.Association -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>

```

```

    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Attribute>
    <Model.ModelElement.name>isDerived</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Feature.visibility xmi.value='public_vis' />
    <Model.Feature.scope xmi.value='instance_level' />
    <Model.StructuralFeature.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>false</XMI.field>
    </Model.StructuralFeature.multiplicity>
    <Model.StructuralFeature.isChangeable xmi.value='true' />
    <Model.Attribute.isDerived xmi.value='false' />
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
    </Model.TypedElement.type>
  </Model.Attribute>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of DataType: AggregationKind _____ -->
<!-- _____ -->
<!-- _____ -->

<Model.DataType xmi.id='a114'>
  <Model.ModelElement.name>AggregationKind</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcEnum xmi.tcName='AggregationKind'
        xmi.tcId='IDL:org.omg.mof/Model/AggregationKind:1.0'>
        <XMI.CorbaTcEnumLabel xmi.tcName='none' />
        <XMI.CorbaTcEnumLabel xmi.tcName='shared' />

```

```

        <XMI.CorbaTcEnumLabel xmi.tcName='composite' />
    </XMI.CorbaTcEnum>
</XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: AssociationEnd -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a90'>
  <Model.ModelElement.name>AssociationEnd</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c40' />
    <Model.Constraint xmi.idref='c41' />
    <Model.Constraint xmi.idref='c42' />
    <Model.Constraint xmi.idref='c43' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c40'>
      <Model.ModelElement.name>EndTypeMustBeClass</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv: self.type.oclIsTypeOf(Class)
        </XMI.any>
      </Model.Constraint.expression>
      <Model.Constraint.language>OCL</Model.Constraint.language>
      <Model.Constraint.evaluationPolicy xmi.value='immediate' />
      <Model.Constraint.constrainedElements>

```

```
<Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c41'>
  <Model.ModelElement.name>EndsMustBeUnique</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
        (self.multiplicity.upper > 1 or
         self.multiplicity.upper = UNBOUNDED) implies
          self.multiplicity.isUnique
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c42'>
  <Model.ModelElement.name>CannotHaveTwoOrderedEnds</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
        self.multiplicity.isOrdered implies
          not self.otherEnd.multiplicity.isOrdered
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c43'>
  <Model.ModelElement.name>CannotHaveTwoAggregateEnds</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
```



```

    inv:
      self.aggregation &lt;&gt; #none implies self.otherEnd = #none
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>isNavigable</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>aggregation</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>

```

```
<Model.DataType xmi.idref='a114' /> <!-- Model.AggregationKind -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>multiplicity</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a99' /> <!-- Model.MultiplicityType -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>isChangeable</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Operation>
  <Model.ModelElement.name>otherEnd</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
```

```

<Model.Feature.scope xmi.value='instance_level' />
<Model.Operation.isQuery xmi.value='true' />
<Model.Namespace.contents>
  <Model.Parameter>
    <Model.ModelElement.name>*return</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Parameter.direction xmi.value='return_dir' />
    <Model.Parameter.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>false</XMI.field>
    </Model.Parameter.multiplicity>
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
    </Model.TypedElement.type>
  </Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of Class: Package _____ -->
<!-- _____ -->
<!-- _____ -->

<Model.Class xmi.id='a50'>
  <Model.ModelElement.name>Package</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c44' />
    <Model.Constraint xmi.idref='c45' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>

```

```

    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
</Model.GeneralizableElement.supertypes>
<Model.Namespace.contents>
  <Model.Constraint xmi.id='c44'>
    <Model.ModelElement.name>PackageContainmentRules</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv:
          Set{Package, Class, DataType, Association, Exception,
            Constraint, Import, Tag}) -&gt; includesAll(self.contentTypes)
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a50' /> <!-- Model.Package -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c45'>
    <Model.ModelElement.name>PackagesCannotBeAbstract</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv: not self.isAbstract
      </XMI.any>
    </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a50' /> <!-- Model.Package -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of DataType: FormatType _____ -->
<!-- _____ -->
<!-- _____ -->

<Model.DataType xmi.id='t2'>

```

```

<Model.ModelElement.name>FormatType</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.DataType.typeCode>
  <XMI.CorbaTypeCode>
    <XMI.CorbaTcAlias xmi.tcName='FormatType'
      xmi.tcId='IDL:org.omg.mof/Model/PackageClass/FormatType:1.0'>
      <XMI.CorbaTypeCode>
        <XMI.CorbaTcString xmi.tcLength='0' />
      </XMI.CorbaTypeCode>
    </XMI.CorbaTcAlias>
  </XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>
<Model.Exception xmi.id='FormatNotSupported'>
  <Model.ModelElement.name>FormatNotSupported</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
</Model.Exception>
<Model.Exception xmi.id='ObjectNotExternalizable'>
  <Model.ModelElement.name>ObjectNotExternalizable</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>explanation</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='out_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a100' /> <!-- Model.string -->

```

```
</Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Exception>
<Model.Exception xmi.id='IllformedExternalizedObject'>
  <Model.ModelElement.name>IllformedExternalizedObject</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>explanation</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='out_dir' />
      <Model.Parameter.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.Parameter.multiplicity>
    <Model.TypedElement.type>
      <Model.DataType xmi.idref='a100' /> <!-- Model.string -->
    </Model.TypedElement.type>
  </Model.Parameter>
</Model.Namespace.contents>
</Model.Exception>
<Model.Operation>
  <Model.ModelElement.name>externalize</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.Operation.isQuery xmi.value='true' />
  <Model.Operation.exceptions>
    <Model.Exception xmi.idref='ObjectNotExternalizable' />
    <Model.Exception xmi.idref='FormatNotSupported' />
  </Model.Operation.exceptions>
  <Model.Namespace.contents>
    <Model.Parameter>
      <Model.ModelElement.name>*return</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Parameter.direction xmi.value='return_dir' />
    </Model.Parameter>
  </Model.Namespace.contents>
</Model.Operation>
</Model.Namespace.contents>
</Model.Exception>
```

```

<Model.Parameter.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a101' /> <!-- Model.any -->
</Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>format</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='t2' /> <!-- Package.FormatType -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
<Model.Operation>
  <Model.ModelElement.name>internalize</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='classifier_level' />
  <Model.Operation.isQuery xmi.value='false' />
  <Model.Operation.exceptions>
    <Model.Exception xmi.idref='FormatNotSupported' />
    <Model.Exception xmi.idref='IllformedExternalizedObject' />
  </Model.Operation.exceptions>
<Model.Namespace.contents>
  <Model.Parameter>
    <Model.ModelElement.name>*return</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Parameter.direction xmi.value='return_dir' />

```

```
<Model.Parameter.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.Parameter.multiplicity>
<Model.TypedElement.type>
  <Model.Class xmi.idref='a50' /> <!-- Model.Package -->
</Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>format</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='t2' /> <!-- Package.FormatType -->
  </Model.TypedElement.type>
</Model.Parameter>
<Model.Parameter>
  <Model.ModelElement.name>stream</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Parameter.direction xmi.value='in_dir' />
  <Model.Parameter.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.Parameter.multiplicity>
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a101' /> <!-- Model.any -->
  </Model.TypedElement.type>
</Model.Parameter>
</Model.Namespace.contents>
</Model.Operation>
</Model.Namespace.contents>
```



```

</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Import -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a68'>
  <Model.ModelElement.name>Import</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c46' />
    <Model.Constraint xmi.idref='c47' />
    <Model.Constraint xmi.idref='c48' />
    <Model.Constraint xmi.idref='c49' />
    <Model.Constraint xmi.idref='c50' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c46'>
      <Model.ModelElement.name>ImportedNamespaceMustBeVisible</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv: self.container.isVisible(self.importedNamespace)
        </XMI.any>
      </Model.Constraint.expression>
      <Model.Constraint.language>OCL</Model.Constraint.language>
      <Model.Constraint.evaluationPolicy xmi.value='deferred' />
      <Model.Constraint.constrainedElements>
        <Model.Class xmi.idref='a68' /> <!-- Model.Import -->
      </Model.Constraint.constrainedElements>
    </Model.Constraint>
  </Model.Namespace.contents>
</Model.Class>

```

```
<Model.Constraint xmi.id='c47'>
  <Model.ModelElement.name>CanOnlyImportPackagesAndClasses</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
        self.imported.oclIsTypeOf(Class) or
        self.imported.oclIsTypeOf(Package)
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a68' /> <!-- Model.Import -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c48'>
  <Model.ModelElement.name>CannotImportSelf</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: self.container &lt;&gt; self.imported
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a68' /> <!-- Model.Import -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c49'>
  <Model.ModelElement.name>CannotImportNestedComponents</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv: not self.container.allContents() -&gt; includes(self.imported)
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
```

```

    <Model.Class xmi.idref='a68' /> <!-- Model.Import -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c50'>
  <Model.ModelElement.name>NestedPackagesCannotImport</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
        self.container -&gt; notEmpty implies
          self.container -&gt; asSequence -&gt; first -&gt; container -&gt; isEmpty
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='deferred' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a68' /> <!-- Model.Import -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>visibility</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a110' /> <!-- Model.VisibilityKind -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>isClustered</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />

```

```

<Model.StructuralFeature.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.Attribute.isDerived xmi.value='false' />
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a98' /> <!-- Model.boolean -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Reference>
  <Model.ModelElement.name>importedNamespace</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a142' /> <!-- Model.Aliases.imported -->
  </Model.Reference.referencedEnd>
</Model.Reference>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of DataType: DirectionKind -->
<!-- _____ -->
<!-- _____ -->

<Model.DataType xmi.id='a112'>

```

```

<Model.ModelElement.name>DirectionKind</Model.ModelElement.name>
<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.DataType.typeCode>
  <XMI.CorbaTypeCode>
    <XMI.CorbaTcEnum xmi.tcName='DirectionKind'
      xmi.tcId='IDL:org.omg.mof/Model/DirectionKind:1.0'>
      <XMI.CorbaTcEnumLabel xmi.tcName='in_dir' />
      <XMI.CorbaTcEnumLabel xmi.tcName='out_dir' />
      <XMI.CorbaTcEnumLabel xmi.tcName='inout_dir' />
      <XMI.CorbaTcEnumLabel xmi.tcName='return_dir' />
    </XMI.CorbaTcEnum>
  </XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Parameter -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a64'>
  <Model.ModelElement.name>Parameter</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Attribute>
      <Model.ModelElement.name>direction</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Feature.visibility xmi.value='public_vis' />

```

```

<Model.Feature.scope xmi.value='instance_level' />
<Model.StructuralFeature.multiplicity>
  <XMI.field>1</XMI.field>
  <XMI.field>1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.Attribute.isDerived xmi.value='false' />
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a112' /> <!-- Model.DirectionKind -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>multiplicity</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a99' /> <!-- Model.MultiplicityType -->
  </Model.TypedElement.type>
</Model.Attribute>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Constraint -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a58'>
  <Model.ModelElement.name>Constraint</Model.ModelElement.name>

```

```

<Model.ModelElement.annotation></Model.ModelElement.annotation>
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='false' />
<Model.GeneralizableElement.isLeaf xmi.value='false' />
<Model.Class.isSingleton xmi.value='false' />
<Model.ModelElement.constraints>
  <Model.Constraint xmi.idref='c51' />
  <Model.Constraint xmi.idref='c52' />
</Model.ModelElement.constraints>
<Model.GeneralizableElement.supertypes>
  <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
</Model.GeneralizableElement.supertypes>
<Model.Namespace.contents>
  <Model.Constraint xmi.id='c51'>
    <Model.ModelElement.name>CannotConstrainThisElement</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv:
          self.constrainedElements -&gt;
          forAll(c |
            not Set{Constraint, Tag, Imports,
              TypeAlias, Constant} -&gt;
            includes(c.oclType())
          </XMI.any>
      </Model.Constraint.expression>
    <Model.Constraint.language>OCL</Model.Constraint.language>
    <Model.Constraint.evaluationPolicy xmi.value='immediate' />
    <Model.Constraint.constrainedElements>
      <Model.Class xmi.idref='a58' /> <!-- Model.Constraint -->
    </Model.Constraint.constrainedElements>
  </Model.Constraint>
  <Model.Constraint xmi.id='c52'>
    <Model.ModelElement.name>ConstraintsLimitedToContainer</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Constraint.expression>
      <XMI.any xmi.type="string" xmi.name="">
        inv:
          self.constrainedElements -&gt;
          forAll(c | self.container.extendedNamespace() -&gt;

```

```
    includes(c))
  </XMI.any>
</Model.Constraint.expression>
<Model.Constraint.language>OCL</Model.Constraint.language>
<Model.Constraint.evaluationPolicy xmi.value='deferred' />
<Model.Constraint.constrainedElements>
  <Model.Class xmi.idref='a58' /> <!-- Model.Constraint -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>expression</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a101' /> <!-- Model.any -->
  </Model.TypedElement.type>
</Model.Attribute>
<Model.Attribute>
  <Model.ModelElement.name>language</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a100' /> <!-- Model.string -->
```



```

    </Model.TypedElement.type>
</Model.Attribute>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of DataType: EvaluationKind -->
<!-- _____ -->
<!-- _____ -->

<Model.DataType xmi.id='a115'>
  <Model.ModelElement.name>EvaluationKind</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcEnum xmi.tcName='EvaluationKind'
        xmi.tcId='IDL:org.omg.mof/Model/NamespaceClass/EvaluationKind:1.0'>
        <XMI.CorbaTcEnumLabel xmi.tcName='immediate' />
        <XMI.CorbaTcEnumLabel xmi.tcName='deferred' />
      </XMI.CorbaTcEnum>
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>
<Model.Attribute>
  <Model.ModelElement.name>evaluationPolicy</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='a115' /> <!-- Constraint.EvaluationKind -->

```

```

    </Model.TypedElement.type>
  </Model.Attribute>
  <Model.Reference>
    <Model.ModelElement.name>constrainedElements</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.Feature.visibility xmi.value='public_vis' />
    <Model.Feature.scope xmi.value='instance_level' />
    <Model.StructuralFeature.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>-1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>true</XMI.field>
    </Model.StructuralFeature.multiplicity>
    <Model.StructuralFeature.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
    </Model.TypedElement.type>
    <Model.Reference.referencedEnd>
      <Model.AssociationEnd xmi.idref='a125' /> <!-- Model.Constrains.constrainedElement -->
    </Model.Reference.referencedEnd>
  </Model.Reference>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: LiteralType -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='t1'>
  <Model.ModelElement.name>LiteralType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcAlias xmi.tcName='LiteralType'
        xmi.tcId='IDL:org.omg.mof/Model/LiteralType:1.0'>

```

```

    <XMI.CorbaTypeCode>
      <XMI.CorbaTcAny/>
    </XMI.CorbaTypeCode>
  </XMI.CorbaTcAlias>
</XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Constant -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a84'>
  <Model.ModelElement.name>Constant</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.ModelElement.constraints>
    <Model.Constraint xmi.idref='c53' />
    <Model.Constraint xmi.idref='c54' />
  </Model.ModelElement.constraints>
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Constraint xmi.id='c53'>
      <Model.ModelElement.name>ConstantsValueMustMatchType</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Constraint.expression>
        <XMI.any xmi.type="string" xmi.name="">
          inv: self.value.type -&gt; equals(self.type -&gt; mapToTypecode())
        </XMI.any>
      </Model.Constraint.expression>
      <Model.Constraint.language>OCL</Model.Constraint.language>
      <Model.Constraint.evaluationPolicy xmi.value='deferred' />
      <Model.Constraint.constrainedElements>

```

```
<Model.Class xmi.idref='a84' /> <!-- Model.Constant -->
</Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Constraint xmi.id='c54'>
  <Model.ModelElement.name>ConstantsTypeMustBeSimpleDataType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Constraint.expression>
    <XMI.any xmi.type="string" xmi.name="">
      inv:
      self.type.oclIsKindOf(DataType) and
      Set{#tk_short, #tk_ushort, #tk_long, #tk_ulong, #tk_char,
      #tk_octet, #tk_float, #tk_double, #tk_boolean, #tk_string,
      #tk_wchar, #tk_wstring, #tk_longlong, #tk_ulonglong,
      #tk_longdouble, #tk_fixed} -&gt;
      includes(self.type.asType(DataType).typecode.
      unwindAliases().kind)
    </XMI.any>
  </Model.Constraint.expression>
  <Model.Constraint.language>OCL</Model.Constraint.language>
  <Model.Constraint.evaluationPolicy xmi.value='immediate' />
  <Model.Constraint.constrainedElements>
    <Model.Class xmi.idref='a84' /> <!-- Model.Constant -->
  </Model.Constraint.constrainedElements>
</Model.Constraint>
<Model.Attribute>
  <Model.ModelElement.name>value</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.Attribute.isDerived xmi.value='false' />
  <Model.TypedElement.type>
    <Model.DataType xmi.idref='t1' /> <!-- Model.LiteralType -->
  </Model.TypedElement.type>
</Model.Attribute>
```

```

    </Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Class: Tag -->
<!-- -->
<!-- _____ -->

<Model.Class xmi.id='a105'>
  <Model.ModelElement.name>Tag</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.Class.isSingleton xmi.value='false' />
  <Model.GeneralizableElement.supertypes>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.GeneralizableElement.supertypes>
  <Model.Namespace.contents>
    <Model.Attribute>
      <Model.ModelElement.name>tagId</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.Feature.visibility xmi.value='public_vis' />
      <Model.Feature.scope xmi.value='instance_level' />
      <Model.StructuralFeature.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.StructuralFeature.multiplicity>
      <Model.StructuralFeature.isChangeable xmi.value='true' />
      <Model.Attribute.isDerived xmi.value='false' />
      <Model.TypedElement.type>
        <Model.DataType xmi.idref='a100' /> <!-- Model.string -->
      </Model.TypedElement.type>
    </Model.Attribute>
    <Model.Attribute>
      <Model.ModelElement.name>values</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>

```

```

<Model.Feature.visibility xmi.value='public_vis' />
<Model.Feature.scope xmi.value='instance_level' />
<Model.StructuralFeature.multiplicity>
  <XMI.field>0</XMI.field>
  <XMI.field>-1</XMI.field>
  <XMI.field>>false</XMI.field>
  <XMI.field>>false</XMI.field>
</Model.StructuralFeature.multiplicity>
<Model.StructuralFeature.isChangeable xmi.value='true' />
<Model.Attribute.isDerived xmi.value='false' />
<Model.TypedElement.type>
  <Model.DataType xmi.idref='a101' /> <!-- Model.any -->
</Model.TypedElement.type>
</Model.Attribute>
<Model.Reference>
  <Model.ModelElement.name>elements</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Feature.visibility xmi.value='public_vis' />
  <Model.Feature.scope xmi.value='instance_level' />
  <Model.StructuralFeature.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.StructuralFeature.multiplicity>
  <Model.StructuralFeature.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
  <Model.Reference.referencedEnd>
    <Model.AssociationEnd xmi.idref='a159' /> <!-- Model.AttachesTo.modelElement -->
  </Model.Reference.referencedEnd>
</Model.Reference>
</Model.Namespace.contents>
</Model.Class>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: AttachesTo -->
<!-- -->

```

```

<!-- _____ -->

<Model.Association xmi.id='a157'>
  <Model.ModelElement.name>AttachesTo</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a159'>
      <Model.ModelElement.name>modelElement</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
    <Model.AssociationEnd xmi.id='a158'>
      <Model.ModelElement.name>tag</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>true</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a105' /> <!-- Model.Tag -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
  </Model.Namespace.contents>
</Model.Association>

```

```

    </Model.TypedElement.type>
  </Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: DependsOn -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a129'>
  <Model.ModelElement.name>DependsOn</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='true' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a130'>
      <Model.ModelElement.name>dependent</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='false' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
    <Model.AssociationEnd xmi.id='a131'>
      <Model.ModelElement.name>provider</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>

```



```

    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.AssociationEnd.multiplicity>
  <Model.AssociationEnd.aggregation xmi.value='none' />
  <Model.AssociationEnd.isNavigable xmi.value='true' />
  <Model.AssociationEnd.isChangeable xmi.value='false' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
</Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: Contains -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a119'>
  <Model.ModelElement.name>Contains</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a120'>
      <Model.ModelElement.name>container</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='composite' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
    </Model.AssociationEnd>
  </Model.Namespace.contents>
</Model.Association>

```

```

    <Model.TypedElement.type>
      <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
<Model.AssociationEnd xmi.id='a121'>
  <Model.ModelElement.name>containedElement</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.AssociationEnd.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>true</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.AssociationEnd.multiplicity>
  <Model.AssociationEnd.aggregation xmi.value='none' />
  <Model.AssociationEnd.isNavigable xmi.value='true' />
  <Model.AssociationEnd.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
  </Model.TypedElement.type>
</Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: Generalizes -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a137'>
  <Model.ModelElement.name>Generalizes</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a138'>
      <Model.ModelElement.name>supertype</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
    </Model.AssociationEnd>
  </Model.Namespace.contents>
</Model.Association>

```

```

<Model.AssociationEnd.multiplicity>
  <XMI.field>0</XMI.field>
  <XMI.field>-1</XMI.field>
  <XMI.field>>true</XMI.field>
  <XMI.field>>true</XMI.field>
</Model.AssociationEnd.multiplicity>
<Model.AssociationEnd.aggregation xmi.value='none' />
<Model.AssociationEnd.isNavigable xmi.value='true' />
<Model.AssociationEnd.isChangeable xmi.value='true' />
<Model.TypedElement.type>
  <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
</Model.TypedElement.type>
</Model.AssociationEnd>
<Model.AssociationEnd xmi.id='a139'>
  <Model.ModelElement.name>subtype</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.AssociationEnd.multiplicity>
    <XMI.field>0</XMI.field>
    <XMI.field>-1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>true</XMI.field>
  </Model.AssociationEnd.multiplicity>
  <Model.AssociationEnd.aggregation xmi.value='none' />
  <Model.AssociationEnd.isNavigable xmi.value='true' />
  <Model.AssociationEnd.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a71' /> <!-- Model.GeneralizableElement -->
  </Model.TypedElement.type>
</Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: Aliases -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a141'>
  <Model.ModelElement.name>Aliases</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>

```

```
<Model.GeneralizableElement.visibility xmi.value='public_vis' />
<Model.GeneralizableElement.isAbstract xmi.value='false' />
<Model.GeneralizableElement.isRoot xmi.value='true' />
<Model.GeneralizableElement.isLeaf xmi.value='true' />
<Model.Association.isDerived xmi.value='false' />
<Model.Namespace.contents>
  <Model.AssociationEnd xmi.id='a143'>
    <Model.ModelElement.name>importer</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.AssociationEnd.multiplicity>
      <XMI.field>0</XMI.field>
      <XMI.field>-1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>true</XMI.field>
    </Model.AssociationEnd.multiplicity>
    <Model.AssociationEnd.aggregation xmi.value='none' />
    <Model.AssociationEnd.isNavigable xmi.value='true' />
    <Model.AssociationEnd.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a68' /> <!-- Model.Import -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
  <Model.AssociationEnd xmi.id='a142'>
    <Model.ModelElement.name>imported</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.AssociationEnd.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>false</XMI.field>
    </Model.AssociationEnd.multiplicity>
    <Model.AssociationEnd.aggregation xmi.value='none' />
    <Model.AssociationEnd.isNavigable xmi.value='true' />
    <Model.AssociationEnd.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a6' /> <!-- Model.Namespace -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>
```

```

<!-- _____ -->
<!-- -->
<!-- Contents of Association: Constrains -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a124'>
  <Model.ModelElement.name>Constrains</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a126'>
      <Model.ModelElement.name>constraint</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a58' /> <!-- Model.Constraint -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
    <Model.AssociationEnd xmi.id='a125'>
      <Model.ModelElement.name>constrainedElement</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />

```

```

    <Model.AssociationEnd.isNavigable xmi.value='true' />
    <Model.AssociationEnd.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a29' /> <!-- Model.ModelElement -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: CanRaise -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a133'>
  <Model.ModelElement.name>CanRaise</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a135'>
      <Model.ModelElement.name>operation</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a25' /> <!-- Model.Operation -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
  </Model.Namespace.contents>
</Model.Association>

```

```

    <Model.ModelElement.name>except</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.AssociationEnd.multiplicity>
      <XMI.field>0</XMI.field>
      <XMI.field>-1</XMI.field>
      <XMI.field>>true</XMI.field>
      <XMI.field>>true</XMI.field>
    </Model.AssociationEnd.multiplicity>
    <Model.AssociationEnd.aggregation xmi.value='none' />
    <Model.AssociationEnd.isNavigable xmi.value='true' />
    <Model.AssociationEnd.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a87' /> <!-- Model.Exception -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- _____ -->
<!-- Contents of Association: Exposes _____ -->
<!-- _____ -->
<!-- _____ -->

<Model.Association xmi.id='a149'>
  <Model.ModelElement.name>Exposes</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='true' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a151'>
      <Model.ModelElement.name>referrer</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd>
    </Model.Namespace.contents>
  </Model.Association>

```

```

</Model.AssociationEnd.multiplicity>
<Model.AssociationEnd.aggregation xmi.value='none' />
<Model.AssociationEnd.isNavigable xmi.value='true' />
<Model.AssociationEnd.isChangeable xmi.value='true' />
<Model.TypedElement.type>
  <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
</Model.TypedElement.type>
</Model.AssociationEnd>
<Model.AssociationEnd xmi.id='a150'>
  <Model.ModelElement.name>exposedEnd</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.AssociationEnd.multiplicity>
    <XMI.field>1</XMI.field>
    <XMI.field>1</XMI.field>
    <XMI.field>>false</XMI.field>
    <XMI.field>>false</XMI.field>
  </Model.AssociationEnd.multiplicity>
  <Model.AssociationEnd.aggregation xmi.value='none' />
  <Model.AssociationEnd.isNavigable xmi.value='true' />
  <Model.AssociationEnd.isChangeable xmi.value='true' />
  <Model.TypedElement.type>
    <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
  </Model.TypedElement.type>
</Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: RefersTo -->
<!-- -->
<!-- _____ -->

<Model.Association xmi.id='a145'>
  <Model.ModelElement.name>RefersTo</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />

```



```

<Model.Namespace.contents>
  <Model.AssociationEnd xmi.id='a147'>
    <Model.ModelElement.name>referent</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.AssociationEnd.multiplicity>
      <XMI.field>0</XMI.field>
      <XMI.field>-1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>true</XMI.field>
    </Model.AssociationEnd.multiplicity>
    <Model.AssociationEnd.aggregation xmi.value='none' />
    <Model.AssociationEnd.isNavigable xmi.value='true' />
    <Model.AssociationEnd.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a86' /> <!-- Model.Reference -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
  <Model.AssociationEnd xmi.id='a146'>
    <Model.ModelElement.name>referencedEnd</Model.ModelElement.name>
    <Model.ModelElement.annotation></Model.ModelElement.annotation>
    <Model.AssociationEnd.multiplicity>
      <XMI.field>1</XMI.field>
      <XMI.field>1</XMI.field>
      <XMI.field>>false</XMI.field>
      <XMI.field>>false</XMI.field>
    </Model.AssociationEnd.multiplicity>
    <Model.AssociationEnd.aggregation xmi.value='none' />
    <Model.AssociationEnd.isNavigable xmi.value='true' />
    <Model.AssociationEnd.isChangeable xmi.value='true' />
    <Model.TypedElement.type>
      <Model.Class xmi.idref='a90' /> <!-- Model.AssociationEnd -->
    </Model.TypedElement.type>
  </Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of Association: IsOfType -->
<!-- -->

```

```
<!-- _____ -->

<Model.Association xmi.id='a153'>
  <Model.ModelElement.name>IsOfType</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='true' />
  <Model.GeneralizableElement.isLeaf xmi.value='true' />
  <Model.Association.isDerived xmi.value='false' />
  <Model.Namespace.contents>
    <Model.AssociationEnd xmi.id='a155'>
      <Model.ModelElement.name>type</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>1</XMI.field>
        <XMI.field>1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>false</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a5' /> <!-- Model.Classifier -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
    <Model.AssociationEnd xmi.id='a154'>
      <Model.ModelElement.name>typedElements</Model.ModelElement.name>
      <Model.ModelElement.annotation></Model.ModelElement.annotation>
      <Model.AssociationEnd.multiplicity>
        <XMI.field>0</XMI.field>
        <XMI.field>-1</XMI.field>
        <XMI.field>>false</XMI.field>
        <XMI.field>>true</XMI.field>
      </Model.AssociationEnd.multiplicity>
      <Model.AssociationEnd.aggregation xmi.value='none' />
      <Model.AssociationEnd.isNavigable xmi.value='true' />
      <Model.AssociationEnd.isChangeable xmi.value='true' />
      <Model.TypedElement.type>
        <Model.Class xmi.idref='a96' /> <!-- Model.TypedElement -->
      </Model.TypedElement.type>
    </Model.AssociationEnd>
  </Model.Namespace.contents>
</Model.Association>
```

```

    </Model.TypedElement.type>
  </Model.AssociationEnd>
</Model.Namespace.contents>
</Model.Association>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: any -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a101'>
  <Model.ModelElement.name>any</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcAny />
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: boolean -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a98'>
  <Model.ModelElement.name>boolean</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcBoolean />
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>

```

```
</XMI.CorbaTypeCode>
</Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: string -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='a100'>
  <Model.ModelElement.name>string</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcString xmi.tcLength='0' />
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- Contents of DataType: unsigned long -->
<!-- -->
<!-- _____ -->

<Model.DataType xmi.id='uLong'>
  <Model.ModelElement.name>unsigned long</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.GeneralizableElement.visibility xmi.value='public_vis' />
  <Model.GeneralizableElement.isAbstract xmi.value='false' />
  <Model.GeneralizableElement.isRoot xmi.value='false' />
  <Model.GeneralizableElement.isLeaf xmi.value='false' />
  <Model.DataType.typeCode>
    <XMI.CorbaTypeCode>
      <XMI.CorbaTcUlong />
    </XMI.CorbaTypeCode>
  </Model.DataType.typeCode>
</Model.DataType>
```

```

    </Model.DataType.typeCode>
</Model.DataType>

<!-- _____ -->
<!-- -->
<!-- IDL Tags -->
<!-- -->
<!-- _____ -->

<Model.Tag>
  <Model.ModelElement.name>IDL Prefix</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Tag.tagId>org.omg.mof.idl_prefix</Model.Tag.tagId>
  <Model.Tag.values>
    <XMI.any xmi.type="string" xmi.name="">org.omg.mof</XMI.any>
  </Model.Tag.values>
  <Model.Tag.elements>
    <Model.Package xmi.idref='a1' /> <!-- Model -->
  </Model.Tag.elements>
</Model.Tag>
<Model.Tag>
  <Model.ModelElement.name>IDL Name for Attribute</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Tag.tagId>org.omg.mof.idl_substitute_name</Model.Tag.tagId>
  <Model.Tag.values>
    <XMI.any xmi.type="string" xmi.name="">MofAttribute</XMI.any>
  </Model.Tag.values>
  <Model.Tag.elements>
    <Model.Class xmi.idref='a23' /> <!-- Model.Attribute -->
  </Model.Tag.elements>
</Model.Tag>
<Model.Tag>
  <Model.ModelElement.name>IDL Name for Exception</Model.ModelElement.name>
  <Model.ModelElement.annotation></Model.ModelElement.annotation>
  <Model.Tag.tagId>org.omg.mof.idl_substitute_name</Model.Tag.tagId>
  <Model.Tag.values>
    <XMI.any xmi.type="string" xmi.name="">MofException</XMI.any>
  </Model.Tag.values>
  <Model.Tag.elements>
    <Model.Class xmi.idref='a87' /> <!-- Model.Exception -->
  </Model.Tag.elements>

```

```
</Model.Tag>

</Model.Namespace.contents>
</Model.Package>
</XMI.content>
</XMI>
```

A.2 The XMI DTD for MOF meta-models

```

<!-- _____ -->
<!-- -->
<!-- XMI is the top-level XML element for XMI transfer text -->
<!-- _____ -->

<!ELEMENT XMI (XMI.header, XMI.content?, XMI.difference*,
              XMI.extensions*) >
<!ATTLIST XMI
          xmi.version CDATA #FIXED "1.0"
          timestamp CDATA #IMPLIED
          verified (true | false) #IMPLIED >

<!-- _____ -->
<!-- -->
<!-- XMI.header contains documentation and identifies the model, -->
<!-- metamodel, and metamodel -->
<!-- _____ -->

<!ELEMENT XMI.header (XMI.documentation?, XMI.model*, XMI.metamodel*,
                    XMI.metamodel*) >

<!-- _____ -->
<!-- -->
<!-- documentation for transfer data -->
<!-- _____ -->

<!ELEMENT XMI.documentation (#PCDATA | XMI.owner | XMI.contact |
                            XMI.longDescription | XMI.shortDescription |
                            XMI.exporter | XMI.exporterVersion |
                            XMI.notice)* >
<!ELEMENT XMI.owner ANY >
<!ELEMENT XMI.contact ANY >
<!ELEMENT XMI.longDescription ANY >
<!ELEMENT XMI.shortDescription ANY >
<!ELEMENT XMI.exporter ANY >
<!ELEMENT XMI.exporterVersion ANY >
<!ELEMENT XMI.exporterID ANY >
<!ELEMENT XMI.notice ANY >

```

```
<!-- _____ -->
<!-- -->
<!-- XMI.element.att defines the attributes that each XML element -->
<!-- that corresponds to a metamodel class must have to conform to -->
<!-- the XMI specification. -->
<!-- _____ -->

<!ENTITY % XMI.element.att
    'xmi.id ID #IMPLIED xmi.label CDATA #IMPLIED xmi.uuid
    CDATA #IMPLIED ' >

<!-- _____ -->
<!-- -->
<!-- XMI.link.att defines the attributes that each XML element that -->
<!-- corresponds to a metamodel class must have to enable it to -->
<!-- function as a simple XLink as well as refer to model -->
<!-- constructs within the same XMI file. -->
<!-- _____ -->

<!ENTITY % XMI.link.att
    'xml:link CDATA #IMPLIED inline (true | false) #IMPLIED
    actuate (show | user) #IMPLIED href CDATA #IMPLIED role
    CDATA #IMPLIED title CDATA #IMPLIED show (embed | replace
    | new) #IMPLIED behavior CDATA #IMPLIED xmi.idref IDREF
    #IMPLIED xmi.uuidref CDATA #IMPLIED' >

<!-- _____ -->
<!-- -->
<!-- XMI.model identifies the model(s) being transferred -->
<!-- _____ -->

<!ELEMENT XMI.model ANY >
<!ATTLIST XMI.model
    %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED >

<!-- _____ -->
<!-- -->
<!-- XMI.metamodel identifies the metamodel(s) for the transferred -->
```

```

<!-- data -->
<!-- _____ -->

<!ELEMENT XMI.metamodel ANY >
<!ATTLIST XMI.metamodel
    %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED >

<!-- _____ -->
<!-- -->
<!-- XMI.metametamodel identifies the metamodel(s) for the -->
<!-- transferred data -->
<!-- _____ -->

<!ELEMENT XMI.metametamodel ANY >
<!ATTLIST XMI.metametamodel
    %XMI.link.att;
    xmi.name CDATA #REQUIRED
    xmi.version CDATA #IMPLIED >

<!-- _____ -->
<!-- -->
<!-- XMI.content is the actual data being transferred -->
<!-- _____ -->

<!ELEMENT XMI.content ANY >

<!-- _____ -->
<!-- -->
<!-- XMI.extensions contains data to transfer that does not conform -->
<!-- to the metamodel(s) in the header -->
<!-- _____ -->

<!ELEMENT XMI.extensions ANY >
<!ATTLIST XMI.extensions
    xmi.extender CDATA #REQUIRED >

<!-- _____ -->
<!-- -->
<!-- extension contains information related to a specific model -->

```

A

```
<!-- construct that is not defined in the metamodel(s) in the header -->
<!-- _____ -->

<!ELEMENT XMI.extension ANY >
<!ATTLIST XMI.extension
    %XMI.element.att;
    %XMI.link.att;
    xmi.extender CDATA #REQUIRED
    xmi.extenderID CDATA #REQUIRED >

<!-- _____ -->
<!-- _____ -->
<!-- XMI.difference holds XML elements representing differences to a -->
<!-- base model -->
<!-- _____ -->

<!ELEMENT XMI.difference (XMI.difference | XMI.delete | XMI.add |
    XMI.replace)* >
<!ATTLIST XMI.difference
    %XMI.element.att;
    %XMI.link.att; >

<!-- _____ -->
<!-- _____ -->
<!-- XMI.delete represents a deletion from a base model -->
<!-- _____ -->

<!ELEMENT XMI.delete EMPTY >
<!ATTLIST XMI.delete
    %XMI.element.att;
    %XMI.link.att; >

<!-- _____ -->
<!-- _____ -->
<!-- XMI.add represents an addition to a base model -->
<!-- _____ -->

<!ELEMENT XMI.add ANY >
<!ATTLIST XMI.add
    %XMI.element.att;
    %XMI.link.att;
```

```

        xmi.position CDATA "-1" >

<!-- _____ -->
<!-- _____ -->
<!-- XMI.replace represents the replacement of a model construct -->
<!-- with another model construct in a base model -->
<!-- _____ -->

<!ELEMENT XMI.replace ANY >
<!ATTLIST XMI.replace
        %XMI.element.att;
        %XMI.link.att;
        xmi.position CDATA "-1" >

<!-- _____ -->
<!-- _____ -->
<!-- XMI.reference may be used to refer to data types not defined in -->
<!-- the metamodel -->
<!-- _____ -->

<!ELEMENT XMI.reference ANY >
<!ATTLIST XMI.reference
        %XMI.link.att; >

<!-- _____ -->
<!-- _____ -->
<!-- This section contains the declaration of XML elements -->
<!-- representing data types -->
<!-- _____ -->

<!ELEMENT XMI.TypeDefinitions ANY >
<!ELEMENT XMI.field ANY >
<!ELEMENT XMI.seqItem ANY >
<!ELEMENT XMI.octetStream (#PCDATA) >
<!ELEMENT XMI.unionDiscrim ANY >

<!ELEMENT XMI.enum EMPTY >
<!ATTLIST XMI.enum xmi.value CDATA #REQUIRED >

<!ELEMENT XMI.any ANY >
<!ATTLIST XMI.any

```

```
%XMI.link.att;
xmi.type CDATA #IMPLIED
xmi.name CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTypeCode (XMI.CorbaTcAlias | XMI.CorbaTcStruct |
                             XMI.CorbaTcSequence | XMI.CorbaTcArray |
                             XMI.CorbaTcEnum | XMI.CorbaTcUnion |
                             XMI.CorbaTcExcept | XMI.CorbaTcString |
                             XMI.CorbaTcWstring | XMI.CorbaTcShort |
                             XMI.CorbaTcLong | XMI.CorbaTcUshort |
                             XMI.CorbaTcUlong | XMI.CorbaTcFloat |
                             XMI.CorbaTcDouble | XMI.CorbaTcBoolean |
                             XMI.CorbaTcChar | XMI.CorbaTcWchar |
                             XMI.CorbaTcOctet | XMI.CorbaTcAny |
                             XMI.CorbaTcTypeCode | XMI.CorbaTcPrincipal |
                             XMI.CorbaTcNull | XMI.CorbaTcVoid |
                             XMI.CorbaTcLongLong | XMI.CorbaTcObjRef |
                             XMI.CorbaTcLongDouble) >

<!ATTLIST XMI.CorbaTypeCode
          %XMI.element.att; >

<!ELEMENT XMI.CorbaTcAlias (XMI.CorbaTypeCode) >
<!ATTLIST XMI.CorbaTcAlias
          xmi.tcName CDATA #REQUIRED
          xmi.tcId   CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTcStruct (XMI.CorbaTcField)* >
<!ATTLIST XMI.CorbaTcStruct
          xmi.tcName CDATA #REQUIRED
          xmi.tcId   CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTcField (XMI.CorbaTypeCode) >
<!ATTLIST XMI.CorbaTcField
          xmi.tcName CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcSequence (XMI.CorbaTypeCode |
                               XMI.CorbaRecursiveType) >
<!ATTLIST XMI.CorbaTcSequence
          xmi.tcLength CDATA #REQUIRED >

<!ELEMENT XMI.CorbaRecursiveType EMPTY >
```

```
<!ATTLIST XMI.CorbaRecursiveType
    xmi.offset CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcArray (XMI.CorbaTypeCode) >
<!ATTLIST XMI.CorbaTcArray
    xmi.tcLength CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcObjRef EMPTY >
<!ATTLIST XMI.CorbaTcObjRef
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTcEnum (XMI.CorbaTcEnumLabel)* >
<!ATTLIST XMI.CorbaTcEnum
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTcEnumLabel EMPTY >
<!ATTLIST XMI.CorbaTcEnumLabel
    xmi.tcName CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcUnionMbr (XMI.CorbaTypeCode, XMI.any) >
<!ATTLIST XMI.CorbaTcUnionMbr
    xmi.tcName CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcUnion (XMI.CorbaTypeCode, XMI.CorbaTcUnionMbr*) >
<!ATTLIST XMI.CorbaTcUnion
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTcExcept (XMI.CorbaTcField)* >
<!ATTLIST XMI.CorbaTcExcept
    xmi.tcName CDATA #REQUIRED
    xmi.tcId CDATA #IMPLIED >

<!ELEMENT XMI.CorbaTcString EMPTY >
<!ATTLIST XMI.CorbaTcString
    xmi.tcLength CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcWstring EMPTY >
<!ATTLIST XMI.CorbaTcWstring
```

```
xmi.tcLength CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcFixed EMPTY >
<!ATTLIST XMI.CorbaTcFixed
    xmi.tcDigits CDATA #REQUIRED
    xmi.tcScale CDATA #REQUIRED >

<!ELEMENT XMI.CorbaTcShort EMPTY >
<!ELEMENT XMI.CorbaTcLong EMPTY >
<!ELEMENT XMI.CorbaTcUshort EMPTY >
<!ELEMENT XMI.CorbaTcUlong EMPTY >
<!ELEMENT XMI.CorbaTcFloat EMPTY >
<!ELEMENT XMI.CorbaTcDouble EMPTY >
<!ELEMENT XMI.CorbaTcBoolean EMPTY >
<!ELEMENT XMI.CorbaTcChar EMPTY >
<!ELEMENT XMI.CorbaTcWchar EMPTY >
<!ELEMENT XMI.CorbaTcOctet EMPTY >
<!ELEMENT XMI.CorbaTcAny EMPTY >
<!ELEMENT XMI.CorbaTcTypeCode EMPTY >
<!ELEMENT XMI.CorbaTcPrincipal EMPTY >
<!ELEMENT XMI.CorbaTcNull EMPTY >
<!ELEMENT XMI.CorbaTcVoid EMPTY >
<!ELEMENT XMI.CorbaTcLongLong EMPTY >
<!ELEMENT XMI.CorbaTcLongDouble EMPTY >

<!-- _____ -->
<!-- -->
<!-- METAMODEL PACKAGE: Model -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.VisibilityKind
    ' xmi.value ( public_vis| protected_vis| private_vis ) #REQUIRED'>

<!ENTITY % Model.DirectionKind
    ' xmi.value ( in_dir| out_dir| inout_dir| return_dir ) #REQUIRED'>

<!ENTITY % Model.ScopeKind
    ' xmi.value ( instance_level| classifier_level ) #REQUIRED'>

<!ENTITY % Model.AggregationKind
```

```

        ' xmi.value ( none| shared| composite ) #REQUIRED'>

<!ELEMENT Model.Namespace.contents ( Model.Feature
|Model.Classifier
|Model.Namespace
|Model.StructuralFeature
|Model.Attribute
|Model.Operation
|Model.Class
|Model.ModelElement
|Model.Package
|Model.Constraint
|Model.Association
|Model.Parameter
|Model.BehavioralFeature
|Model.Import
|Model.GeneralizableElement
|Model.Constant
|Model.Reference
|Model.Exception
|Model.DataType
|Model.AssociationEnd
|Model.TypedElement
|Model.TypeAlias
|Model.Tag)* >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.ModelElement -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.ModelElement.DepthKind
        ' xmi.value ( shallow| deep ) #REQUIRED'>

<!ENTITY % Model.ModelElement.VerifyResultKind
        ' xmi.value ( valid| invalid| published ) #REQUIRED'>

<!ELEMENT Model.ModelElement.name (#PCDATA|XMI.reference)*>

```

```

<!ELEMENT Model.ModelElement.annotation (#PCDATA|XMI.reference)*>

<!ELEMENT Model.ModelElement.container (Model.Classifier
|Model.Namespace
|Model.Operation
|Model.Class
|Model.Package
|Model.Association
|Model.BehavioralFeature
|Model.GeneralizableElement
|Model.Exception
|Model.DataType)?>

<!ELEMENT Model.ModelElement.constraints (Model.Constraint)*>

<!ENTITY % Model.ModelElementProperties '((Model.ModelElement.name)?
,(Model.ModelElement.annotation?))' >

<!ENTITY % Model.ModelElementAssociations '(Model.ModelElement.container?
,Model.ModelElement.constraints*)' >

<!ELEMENT Model.ModelElement ( %Model.ModelElementProperties;
,(XMI.extension* , %Model.ModelElementAssociations; ) )?>

<!ATTLIST Model.ModelElement %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Feature -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Feature.visibility EMPTY>
<!ATTLIST Model.Feature.visibility %Model.VisibilityKind;>

<!ELEMENT Model.Feature.scope EMPTY>
<!ATTLIST Model.Feature.scope %Model.ScopeKind;>

<!ENTITY % Model.FeatureProperties '(%Model.ModelElementProperties;
,(Model.Feature.visibility)?

```



```

    ,(Model.Feature.scope)?)' >

<!ENTITY % Model.FeatureAssociations '(%Model.ModelElementAssociations;)' >

<!ELEMENT Model.Feature ( %Model.FeatureProperties;
    ,(XMI.extension* , %Model.FeatureAssociations; ) )?>

<!ATTLIST Model.Feature %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Namespace -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.NamespaceProperties '(%Model.ModelElementProperties;)' >

<!ENTITY % Model.NamespaceAssociations '(%Model.ModelElementAssociations;)' >

<!ENTITY % Model.NamespaceCompositions '(Model.Namespace.contents*)' >

<!ELEMENT Model.Namespace ( %Model.NamespaceProperties;
    ,(XMI.extension* , %Model.NamespaceAssociations; )
    , %Model.NamespaceCompositions; )?>

<!ATTLIST Model.Namespace %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.GeneralizableElement -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.GeneralizableElement.visibility EMPTY>
<!ATTLIST Model.GeneralizableElement.visibility %Model.VisibilityKind;>

<!ELEMENT Model.GeneralizableElement.isAbstract EMPTY>
<!ATTLIST Model.GeneralizableElement.isAbstract
    xmi.value ( true | false ) #REQUIRED>

```

```

<!ELEMENT Model.GeneralizableElement.isRoot EMPTY>
<!ATTLIST Model.GeneralizableElement.isRoot
    xmi.value ( true | false ) #REQUIRED>

<!ELEMENT Model.GeneralizableElement.isLeaf EMPTY>
<!ATTLIST Model.GeneralizableElement.isLeaf
    xmi.value ( true | false ) #REQUIRED>

<!ELEMENT Model.GeneralizableElement.supertypes (Model.Classifier
|Model.Class
|Model.Package
|Model.Association
|Model.GeneralizableElement
|Model.DataType)*>

<!ENTITY % Model.GeneralizableElementProperties '(%Model.NamespaceProperties;
, (Model.GeneralizableElement.visibility)?
, (Model.GeneralizableElement.isAbstract)?
, (Model.GeneralizableElement.isRoot)?
, (Model.GeneralizableElement.isLeaf)?)' >

<!ENTITY % Model.GeneralizableElementAssociations '(%Model.NamespaceAssociations;
, Model.GeneralizableElement.supertypes*)' >

<!ENTITY % Model.GeneralizableElementCompositions '(%Model.NamespaceCompositions;)' >

<!ELEMENT Model.GeneralizableElement ( %Model.GeneralizableElementProperties;
, (XMI.extension* , %Model.GeneralizableElementAssociations; )
, %Model.GeneralizableElementCompositions; )?>

<!ATTLIST Model.GeneralizableElement %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Classifier -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.ClassifierProperties '(%Model.GeneralizableElementProperties;)' >

```

```

<!ENTITY % Model.ClassifierAssociations '(%Model.GeneralizableElementAssociations;)' >

<!ENTITY % Model.ClassifierCompositions '(%Model.GeneralizableElementCompositions;)' >

<!ELEMENT Model.Classifier ( %Model.ClassifierProperties;
    ,(XMI.extension* , %Model.ClassifierAssociations; )
    , %Model.ClassifierCompositions; )?>

<!ATTLIST Model.Classifier %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.TypedElement -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.TypedElement.type (Model.Classifier
|Model.Class
|Model.Association
|Model.DataType)?>

<!ENTITY % Model.TypedElementProperties '(%Model.ModelElementProperties;)' >

<!ENTITY % Model.TypedElementAssociations '(%Model.ModelElementAssociations;
,Model.TypedElement.type?)' >

<!ELEMENT Model.TypedElement ( %Model.TypedElementProperties;
    ,(XMI.extension* , %Model.TypedElementAssociations; ) )?>

<!ATTLIST Model.TypedElement %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.StructuralFeature -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.StructuralFeature.multiplicity (XMI.field|XMI.reference)*>

```

```
<!ELEMENT Model.StructuralFeature.isChangeable EMPTY>
<!ATTLIST Model.StructuralFeature.isChangeable
    xmi.value ( true | false ) #REQUIRED>

<!ENTITY % Model.StructuralFeatureProperties '(%Model.FeatureProperties;
    ,(Model.StructuralFeature.multiplicity)?
    ,(Model.StructuralFeature.isChangeable)?)' >

<!ENTITY % Model.StructuralFeatureAssociations '(%Model.FeatureAssociations;
    ,Model.TypedElement.type?)' >

<!ELEMENT Model.StructuralFeature ( %Model.StructuralFeatureProperties;
    ,(XMI.extension* , %Model.StructuralFeatureAssociations; ) )?>

<!ATTLIST Model.StructuralFeature %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- _____ -->
<!-- METAMODEL CLASS: Model.Attribute -->
<!-- _____ -->
<!-- _____ -->

<!ELEMENT Model.Attribute.isDerived EMPTY>
<!ATTLIST Model.Attribute.isDerived
    xmi.value ( true | false ) #REQUIRED>

<!ENTITY % Model.AttributeProperties '(%Model.StructuralFeatureProperties;
    ,(Model.Attribute.isDerived)?)' >

<!ENTITY % Model.AttributeAssociations '(%Model.StructuralFeatureAssociations;)' >

<!ELEMENT Model.Attribute ( %Model.AttributeProperties;
    ,(XMI.extension* , %Model.AttributeAssociations; ) )?>

<!ATTLIST Model.Attribute %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- _____ -->
```

```

<!-- METAMODEL CLASS: Model.BehavioralFeature -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.BehavioralFeatureProperties '(%Model.FeatureProperties;)' >

<!ENTITY % Model.BehavioralFeatureAssociations '(%Model.FeatureAssociations;)' >

<!ENTITY % Model.BehavioralFeatureCompositions '(Model.Namespace.contents*)' >

<!ELEMENT Model.BehavioralFeature ( %Model.BehavioralFeatureProperties;
    ,(XMI.extension* , %Model.BehavioralFeatureAssociations; )
    , %Model.BehavioralFeatureCompositions; )?>

<!ATTLIST Model.BehavioralFeature %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Operation -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Operation.isQuery EMPTY>
<!ATTLIST Model.Operation.isQuery
    xmi.value ( true | false ) #REQUIRED>

<!ELEMENT Model.Operation.exceptions (Model.Exception)*>

<!ENTITY % Model.OperationProperties '(%Model.BehavioralFeatureProperties;
    ,(Model.Operation.isQuery?))' >

<!ENTITY % Model.OperationAssociations '(%Model.BehavioralFeatureAssociations;
    ,Model.Operation.exceptions*)' >

<!ENTITY % Model.OperationCompositions '(%Model.BehavioralFeatureCompositions;)' >

<!ELEMENT Model.Operation ( %Model.OperationProperties;
    ,(XMI.extension* , %Model.OperationAssociations; )
    , %Model.OperationCompositions; )?>

```

```
<!ATTLIST Model.Operation %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Class -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Class.isSingleton EMPTY>
<!ATTLIST Model.Class.isSingleton
    xmi.value ( true | false ) #REQUIRED>

<!ENTITY % Model.ClassProperties '(%Model.ClassifierProperties;
    ,(Model.Class.isSingleton?))' >

<!ENTITY % Model.ClassAssociations '(%Model.ClassifierAssociations;)' >

<!ENTITY % Model.ClassCompositions '(%Model.ClassifierCompositions;)' >

<!ELEMENT Model.Class ( %Model.ClassProperties;
    ,(XMI.extension* , %Model.ClassAssociations; )
    , %Model.ClassCompositions; )?>

<!ATTLIST Model.Class %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Package -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.PackageProperties '(%Model.GeneralizableElementProperties;)' >

<!ENTITY % Model.PackageAssociations '(%Model.GeneralizableElementAssociations;)' >

<!ENTITY % Model.PackageCompositions '(%Model.GeneralizableElementCompositions;)' >

<!ELEMENT Model.Package ( %Model.PackageProperties;
    ,(XMI.extension* , %Model.PackageAssociations; )
```

```

    , %Model.PackageCompositions; )?>

<!ATTLIST Model.Package %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Constraint -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.Constraint.EvaluationKind
    ' xmi.value ( immediate| deferred ) #REQUIRED'>

<!ELEMENT Model.Constraint.expression (XMI.any)>

<!ELEMENT Model.Constraint.language (#PCDATA|XMI.reference)*>

<!ELEMENT Model.Constraint.evaluationPolicy EMPTY>
<!ATTLIST Model.Constraint.evaluationPolicy %Model.Constraint.EvaluationKind;>

<!ELEMENT Model.Constraint.constrainedElements (Model.Feature
|Model.Classifier
|Model.Namespace
|Model.StructuralFeature
|Model.Attribute
|Model.Operation
|Model.Class
|Model.ModelElement
|Model.Package
|Model.Constraint
|Model.Association
|Model.Parameter
|Model.BehavioralFeature
|Model.Import
|Model.GeneralizableElement
|Model.Constant
|Model.Reference
|Model.Exception
|Model.DataType
|Model.AssociationEnd

```

```
|Model.TypedElement
|Model.TypeAlias
|Model.Tag)*>

<!ENTITY % Model.ConstraintProperties '(%Model.ModelElementProperties;
, (Model.Constraint.expression)?
, (Model.Constraint.language)?
, (Model.Constraint.evaluationPolicy)?)' >

<!ENTITY % Model.ConstraintAssociations '(%Model.ModelElementAssociations;
, Model.Constraint.constrainedElements*)' >

<!ELEMENT Model.Constraint ( %Model.ConstraintProperties;
, (XMI.extension* , %Model.ConstraintAssociations; ) )?>

<!ATTLIST Model.Constraint %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Association -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Association.isDerived EMPTY>
<!ATTLIST Model.Association.isDerived
xmi.value ( true | false ) #REQUIRED>

<!ENTITY % Model.AssociationProperties '(%Model.ClassifierProperties;
, (Model.Association.isDerived)?)' >

<!ENTITY % Model.AssociationAssociations '(%Model.ClassifierAssociations;)' >

<!ENTITY % Model.AssociationCompositions '(%Model.ClassifierCompositions;)' >

<!ELEMENT Model.Association ( %Model.AssociationProperties;
, (XMI.extension* , %Model.AssociationAssociations; )
, %Model.AssociationCompositions; )?>

<!ATTLIST Model.Association %XMI.element.att; %XMI.link.att; >
```



```

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Parameter -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Parameter.direction EMPTY>
<!ATTLIST Model.Parameter.direction %Model.DirectionKind;>

<!ELEMENT Model.Parameter.multiplicity (XMI.field|XMI.reference)*>

<!ENTITY % Model.ParameterProperties '(%Model.TypedElementProperties;
, (Model.Parameter.direction)?
, (Model.Parameter.multiplicity)?)' >

<!ENTITY % Model.ParameterAssociations '(%Model.TypedElementAssociations;)' >

<!ELEMENT Model.Parameter ( %Model.ParameterProperties;
, (XMI.extension* , %Model.ParameterAssociations; ) )?>

<!ATTLIST Model.Parameter %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Import -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Import.visibility EMPTY>
<!ATTLIST Model.Import.visibility %Model.VisibilityKind;>

<!ELEMENT Model.Import.isClustered EMPTY>
<!ATTLIST Model.Import.isClustered
xmi.value ( true | false ) #REQUIRED>

<!ELEMENT Model.Import.importedNamespace (Model.Classifier
|Model.Namespace
|Model.Operation
|Model.Class

```

```
|Model.Package
|Model.Association
|Model.BehavioralFeature
|Model.GeneralizableElement
|Model.Exception
|Model.DataType)?>

<!ENTITY % Model.ImportProperties '(%Model.ModelElementProperties;
, (Model.Import.visibility)?
, (Model.Import.isClustered)?)' >

<!ENTITY % Model.ImportAssociations '(%Model.ModelElementAssociations;
, Model.Import.importedNamespace?)' >

<!ELEMENT Model.Import ( %Model.ImportProperties;
, (XMI.extension* , %Model.ImportAssociations; ) )?>

<!ATTLIST Model.Import %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Constant -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Constant.value (XMI.any)>

<!ENTITY % Model.ConstantProperties '(%Model.TypedElementProperties;
, (Model.Constant.value)?)' >

<!ENTITY % Model.ConstantAssociations '(%Model.TypedElementAssociations;)' >

<!ELEMENT Model.Constant ( %Model.ConstantProperties;
, (XMI.extension* , %Model.ConstantAssociations; ) )?>

<!ATTLIST Model.Constant %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
```

```

<!-- METAMODEL CLASS: Model.Reference -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Reference.referencedEnd (Model.AssociationEnd)?>

<!ENTITY % Model.ReferenceProperties '(%Model.StructuralFeatureProperties;)' >

<!ENTITY % Model.ReferenceAssociations '(%Model.StructuralFeatureAssociations;
,Model.Reference.referencedEnd?)' >

<!ELEMENT Model.Reference ( %Model.ReferenceProperties;
,(XMI.extension* , %Model.ReferenceAssociations; ) )?>

<!ATTLIST Model.Reference %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Exception -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.ExceptionProperties '(%Model.BehavioralFeatureProperties;)' >

<!ENTITY % Model.ExceptionAssociations '(%Model.BehavioralFeatureAssociations;)' >

<!ENTITY % Model.ExceptionCompositions '(%Model.BehavioralFeatureCompositions;)' >

<!ELEMENT Model.Exception ( %Model.ExceptionProperties;
,(XMI.extension* , %Model.ExceptionAssociations; )
, %Model.ExceptionCompositions; )?>

<!ATTLIST Model.Exception %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.DataType -->
<!-- -->
<!-- _____ -->

```

```
<!ELEMENT Model.DataType.typeCode (XMI.CorbaTypeCode)>

<!ENTITY % Model.DataTypeProperties '(%Model.ClassifierProperties;
    ,(Model.DataType.typeCode?))' >

<!ENTITY % Model.DataTypeAssociations '(%Model.ClassifierAssociations;)' >

<!ENTITY % Model.DataTypeCompositions '(%Model.ClassifierCompositions;)' >

<!ELEMENT Model.DataType ( %Model.DataTypeProperties;
    ,(XMI.extension* , %Model.DataTypeAssociations; )
    , %Model.DataTypeCompositions; )?>

<!ATTLIST Model.DataType %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.AssociationEnd -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.AssociationEnd.multiplicity (XMI.field|XMI.reference)*>

<!ELEMENT Model.AssociationEnd.aggregation EMPTY>
<!ATTLIST Model.AssociationEnd.aggregation %Model.AggregationKind;>

<!ELEMENT Model.AssociationEnd.isNavigable EMPTY>
<!ATTLIST Model.AssociationEnd.isNavigable
    xmi.value ( true | false ) #REQUIRED>

<!ELEMENT Model.AssociationEnd.isChangeable EMPTY>
<!ATTLIST Model.AssociationEnd.isChangeable
    xmi.value ( true | false ) #REQUIRED>

<!ENTITY % Model.AssociationEndProperties '(%Model.TypedElementProperties;
    ,(Model.AssociationEnd.multiplicity)?
    ,(Model.AssociationEnd.aggregation)?
    ,(Model.AssociationEnd.isNavigable)?
    ,(Model.AssociationEnd.isChangeable?))' >
```

```

<!ENTITY % Model.AssociationEndAssociations '(%Model.TypedElementAssociations;)' >

<!ELEMENT Model.AssociationEnd ( %Model.AssociationEndProperties;
    ,(XMI.extension* , %Model.AssociationEndAssociations; ) )?>

<!ATTLIST Model.AssociationEnd %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.TypeAlias -->
<!-- -->
<!-- _____ -->

<!ENTITY % Model.TypeAliasProperties '(%Model.TypedElementProperties;)' >

<!ENTITY % Model.TypeAliasAssociations '(%Model.TypedElementAssociations;)' >

<!ELEMENT Model.TypeAlias ( %Model.TypeAliasProperties;
    ,(XMI.extension* , %Model.TypeAliasAssociations; ) )?>

<!ATTLIST Model.TypeAlias %XMI.element.att; %XMI.link.att; >

<!-- _____ -->
<!-- -->
<!-- METAMODEL CLASS: Model.Tag -->
<!-- -->
<!-- _____ -->

<!ELEMENT Model.Tag.tagId (#PCDATA|XMI.reference)*>

<!ELEMENT Model.Tag.values (XMI.any)>

<!ELEMENT Model.Tag.elements (Model.Feature
    |Model.Classifier
    |Model.Namespace
    |Model.StructuralFeature
    |Model.Attribute
    |Model.Operation

```

```
|Model.Class
|Model.ModelElement
|Model.Package
|Model.Constraint
|Model.Association
|Model.Parameter
|Model.BehavioralFeature
|Model.Import
|Model.GeneralizableElement
|Model.Constant
|Model.Reference
|Model.Exception
|Model.DataType
|Model.AssociationEnd
|Model.TypedElement
|Model.TypeAlias
|Model.Tag)*>

<!ENTITY % Model.TagProperties '(%Model.ModelElementProperties;
, (Model.Tag.tagId)?
, (Model.Tag.values)*)' >

<!ENTITY % Model.TagAssociations '(%Model.ModelElementAssociations;
, Model.Tag.elements*)' >

<!ELEMENT Model.Tag ( %Model.TagProperties;
, (XMI.extension* , %Model.TagAssociations; ) )?>

<!ATTLIST Model.Tag %XMI.element.att; %XMI.link.att; >

<!ELEMENT Model ((Model.Feature
|Model.Classifier
|Model.Namespace
|Model.StructuralFeature
|Model.Attribute
|Model.Operation
|Model.Class
|Model.ModelElement
|Model.Package
|Model.Constraint
|Model.Association
```

```
|Model.Parameter  
|Model.BehavioralFeature  
|Model.Import  
|Model.GeneralizableElement  
|Model.Constant  
|Model.Reference  
|Model.Exception  
|Model.DataType  
|Model.AssociationEnd  
|Model.TypedElement  
|Model.TypeAlias  
|Model.Tag)* )>  
<!ATTLIST Model %XMI.element.att; %XMI.link.att;>
```


MOF IDL Summary

B

This appendix summarizes the CORBA IDL for the Meta Object Facility so that it can be easily processed by IDL compilers. The IDL for the Model and Reflective packages has been included so that the appropriate sections can be compiled separately as needed.

B.1 MOF Model IDL

```
#pragma prefix "org.omg.mof"
module Model {
    typedef sequence < any > AnyBag;
    interface ModelPackage;
    interface ModelElementClass;
    interface ModelElement;
    typedef sequence <::Model::ModelElement> ModelElementSet;
    typedef sequence <::Model::ModelElement> ModelElementUList;
    interface NamespaceClass;
    interface Namespace;
    typedef sequence <::Model::Namespace> NamespaceUList;
    interface GeneralizableElementClass;
    interface GeneralizableElement;
    typedef sequence <::Model::GeneralizableElement> GeneralizableElementSet;
    typedef sequence <::Model::GeneralizableElement> GeneralizableElementUList;
    interface TypedElementClass;
    interface TypedElement;
    typedef sequence <::Model::TypedElement> TypedElementSet;
    typedef sequence <::Model::TypedElement> TypedElementUList;
    interface ClassifierClass;
    interface Classifier;
    typedef sequence <::Model::Classifier> ClassifierUList;
    interface ClassClass;
    interface Class;
    typedef sequence <::Model::Class> ClassUList;
    interface DataTypeClass;
    interface DataType;
```

```
typedef sequence <::Model::DataType> DataTypeUList;
interface TypeAliasClass;
interface TypeAlias;
typedef sequence <::Model::TypeAlias> TypeAliasUList;
interface FeatureClass;
interface Feature;
typedef sequence <::Model::Feature> FeatureUList;
interface StructuralFeatureClass;
interface StructuralFeature;
typedef sequence <::Model::StructuralFeature> StructuralFeatureUList;
interface MofAttributeClass;
interface MofAttribute;
typedef sequence <::Model::MofAttribute> MofAttributeUList;
interface ReferenceClass;
interface Reference;
typedef sequence <::Model::Reference> ReferenceSet;
typedef sequence <::Model::Reference> ReferenceUList;
interface BehavioralFeatureClass;
interface BehavioralFeature;
typedef sequence <::Model::BehavioralFeature> BehavioralFeatureUList;
interface OperationClass;
interface Operation;
typedef sequence <::Model::Operation> OperationSet;
typedef sequence <::Model::Operation> OperationUList;
interface MofExceptionClass;
interface MofException;
typedef sequence <::Model::MofException> MofExceptionUList;
interface AssociationClass;
interface Association;
typedef sequence <::Model::Association> AssociationUList;
interface AssociationEndClass;
interface AssociationEnd;
typedef sequence <::Model::AssociationEnd> AssociationEndUList;
interface PackageClass;
interface Package;
typedef sequence <::Model::Package> PackageUList;
interface ImportClass;
interface Import;
typedef sequence <::Model::Import> ImportSet;
typedef sequence <::Model::Import> ImportUList;
interface ParameterClass;
interface Parameter;
typedef sequence <::Model::Parameter> ParameterUList;
interface ConstraintClass;
interface Constraint;
typedef sequence <::Model::Constraint> ConstraintSet;
typedef sequence <::Model::Constraint> ConstraintUList;
interface ConstantClass;
interface Constant;
typedef sequence <::Model::Constant> ConstantUList;
interface TagClass;
interface Tag;
typedef sequence <::Model::Tag> TagUList;
typedef string NameType;
typedef sequence <::Model::NameType> NameTypeList;
```

```

typedef string AnnotationType;
interface ModelElementClass : Reflective::RefObject {
    readonly attribute ModelElementUList all_of_type_model_element;
    const string MUST_BE_CONTAINED_UNLESS_PACKAGE =
        "org.omg.mof:constraint.model.model_element.must_be_contained_unless_package";
    const string FROZEN_ATTRIBUTES_CANNOT_BE_CHANGED =
        "org.omg.mof:constraint.model.model_element.frozen_attributes_cannot_be_changed";
    const string FROZEN_ELEMENTS_CANNOT_BE_DELETED =
        "org.omg.mof:constraint.model.model_element.frozen_elements_cannot_be_deleted";
    const string FROZEN_DEPENDENCIES_CANNOT_BE_CHANGED =
        "org.omg.mof:constraint.model.model_element.frozen_dependencies_cannot_be_changed";
    typedef string DependencyKind;
    typedef sequence <::Model::ModelElementClass::DependencyKind> DependencyKindSet;
    const DependencyKind CONTAINER_DEP = "container";
    const DependencyKind CONTENTS_DEP = "contents";
    const DependencyKind SIGNATURE_DEP = "signature";
    const DependencyKind CONSTRAINT_DEP = "constraint";
    const DependencyKind CONSTRAINED_ELEMENTS_DEP = "constrained elements";
    const DependencyKind SPECIALIZATION_DEP = "specialization";
    const DependencyKind IMPORT_DEP = "import";
    const DependencyKind TYPE_DEFINITION_DEP = "type definition";
    const DependencyKind REFERENCED_ENDS_DEP = "referenced ends";
    const DependencyKind TAGGED_ELEMENTS_DEP = "tagged elements";
    const DependencyKind INDIRECT_DEP = "indirect";
    const DependencyKind ALL_DEP = "all";
    enum VerifyResultKind { valid, published, invalid };
    enum DepthKind { shallow, deep };
    struct ViolationType {
        string error_kind;
        Reflective::RefObject element_in_error;
        Reflective::NamedValueList values_in_error;
        string error_description;
    };
    typedef sequence <::Model::ModelElementClass::ViolationType> ViolationTypeSet;
}; // end of interface ModelElementClass

interface ModelElement : ModelElementClass {
    NameType name ()
        raises (Reflective::MofError);
    void set_name (in NameType new_value)
        raises (Reflective::MofError);
    NameTypeList qualified_name ()
        raises (Reflective::MofError);
    AnnotationType annotation ()
        raises (Reflective::MofError);
    void set_annotation (in AnnotationType new_value)
        raises (Reflective::MofError);
    ModelElementSet required_elements ()
        raises (Reflective::MofError);
    ModelElementSet find_required_elements (
        in ModelElementClass::DependencyKindSet kinds,
        in boolean recursive)
        raises (Reflective::MofError);
    boolean is_required_because (in ModelElement other_element,

```

```

        out ModelElementClass::DependencyKind reason)
    raises (Reflective::MofError);
Namespace container ()
    raises (Reflective::NotSet, Reflective::MofError);
void set_container (in Namespace new_value)
    raises (Reflective::MofError);
void unset_container ()
    raises (Reflective::MofError);
ConstraintSet constraints ()
    raises (Reflective::MofError);
void set_constraints (in ConstraintSet new_value)
    raises (Reflective::MofError);
void add_constraints (in Constraint new_element)
    raises (Reflective::MofError);
void modify_constraints (in Constraint old_element, in Constraint new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_constraints (in Constraint old_element)
    raises (Reflective::NotFound, Reflective::MofError);
ModelElementClass::VerifyResultKind verify (
    in ModelElementClass::DepthKind depth,
    out ModelElementClass::ViolationTypeSet problems)
    raises (Reflective::MofError);
boolean is_frozen ()
    raises (Reflective::MofError);
boolean is_visible (in ModelElement other_element)
    raises (Reflective::MofError);
}; // end of interface ModelElement

enum VisibilityKind { public_vis, protected_vis, private_vis };

interface NamespaceClass : ModelElementClass {
    readonly attribute NamespaceUList all_of_type_namespace;
    const string CONTENT_NAMES_MUST_NOT_COLLIDE =
        "org.omg.mof.constraint.model.namespace.content_names_must_not_collide";
    exception NameNotFound { NameType name; };
    exception NameNotResolved {
        string explanation;
        NameTypeList rest_of_name; };
}; // end of interface NamespaceClass

interface Namespace : NamespaceClass, ModelElement {
    ModelElementUList contents ()
        raises (Reflective::MofError);
    void set_contents (in ModelElementUList new_value)
        raises (Reflective::MofError);
    void add_contents (in ModelElement new_element)
        raises (Reflective::MofError);
    void add_contents_before (in ModelElement new_element, in ModelElement before_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_contents (in ModelElement old_element, in ModelElement new_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove_contents (in ModelElement old_element)
        raises (Reflective::NotFound, Reflective::MofError);
    ModelElement lookup_element (in NameType name)
        raises (NamespaceClass::NameNotFound, Reflective::MofError);
};

```

```

    ModelElement resolve_qualified_name (in NameTypeList qualified_name)
        raises (NamespaceClass::NameNotResolved, Reflective::MofError);
    ModelElementUList find_elements_by_type (in Class of_type, in boolean include_subtypes)
        raises (Reflective::MofError);
    boolean name_is_valid (in NameType proposed_name)
        raises (Reflective::MofError);
}; // end of interface Namespace

interface GeneralizableElementClass : NamespaceClass {
    readonly attribute GeneralizableElementUList
        all_of_type_generalizable_element;
    const string SUPERTYPE_MUST_NOT_BE_SELF =
        "org.omg.mof:constraint.model.generalizable_element.supertype_must_not_be_self";
    const string SUPERTYPE_KIND_MUST_BE_SAME =
        "org.omg.mof:constraint.model.generalizable_element.supertype_kind_must_be_same";
    const string CONTENTS_MUST_NOT_COLLIDE_WITH_SUPERTYPES =
        "org.omg.mof:constraint.model.generalizable_element"
        ".contents_must_not_collide_with_supertypes";
    const string DIAMOND_RULE_MUST_BE_OBEYED =
        "org.omg.mof:constraint.model.generalizable_element.diamond_rule_must_be_obeyed";
    const string NO_SUPERTYPES_ALLOWED_FOR_ROOT =
        "org.omg.mof:constraint.model.generalizable_element.no_supertypes_allowed_for_root";
    const string SUPERTYPES_MUST_BE_VISIBLE =
        "org.omg.mof:constraint.model.generalizable_element.supertypes_must_be_visible";
    const string NO_SUBTYPES_ALLOWED_FOR_LEAF =
        "org.omg.mof:constraint.model.generalizable_element.no_subtypes_allowed_for_leaf";
}; // end of interface GeneralizableElementClass

interface GeneralizableElement : GeneralizableElementClass, Namespace {
    boolean is_root ()
        raises (Reflective::MofError);
    void set_is_root (in boolean new_value)
        raises (Reflective::MofError);
    boolean is_leaf ()
        raises (Reflective::MofError);
    void set_is_leaf (in boolean new_value)
        raises (Reflective::MofError);
    boolean is_abstract ()
        raises (Reflective::MofError);
    void set_is_abstract (in boolean new_value)
        raises (Reflective::MofError);
    VisibilityKind visibility ()
        raises (Reflective::MofError);
    void set_visibility (in VisibilityKind new_value)
        raises (Reflective::MofError);
    GeneralizableElementUList supertypes ()
        raises (Reflective::MofError);
    void set_supertypes (in GeneralizableElementUList new_value)
        raises (Reflective::MofError);
    void add_supertypes (in GeneralizableElement new_element)
        raises (Reflective::MofError);
    void add_supertypes_before (in GeneralizableElement new_element,
                               in GeneralizableElement before_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_supertypes (in GeneralizableElement old_element,

```

```

        in GeneralizableElement new_element)
    raises (Reflective::NotFound, Reflective::MofError);
void remove_supertypes (in GeneralizableElement old_element)
    raises (Reflective::NotFound, Reflective::MofError);
GeneralizableElementSet all_supertypes ()
    raises (Reflective::MofError);
ModelElement lookup_element_extended (in NameType name)
    raises (Reflective::MofError);
ModelElementUList find_elements_by_type_extended (
    in Class of_type,
    in boolean include_subtypes)
    raises (Reflective::MofError);
}; // end of interface GeneralizableElement

interface TypedElementClass : ModelElementClass {
    // get all typed_element including subtypes of typed_element
    readonly attribute TypedElementUList all_of_type_typed_element;
    const string ASSOCIATIONS_CANNOT_BE_TYPES =
        "org.omg.mof:constraint.model.typed_element.associations_cannot_be_types";
    const string TYPE_MUST_BE_VISIBLE =
        "org.omg.mof:constraint.model.typed_element.type_must_be_visible";
}; // end of interface TypedElementClass

interface TypedElement : TypedElementClass, ModelElement {
    Classifier type ()
        raises (Reflective::MofError);
    void set_type (in Classifier new_value)
        raises (Reflective::MofError);
}; // end of interface TypedElement

interface ClassifierClass : GeneralizableElementClass {
    readonly attribute ClassifierUList all_of_type_classifier;
}; // end of interface ClassifierClass

interface Classifier : ClassifierClass, GeneralizableElement { };

interface ClassClass : ClassifierClass {
    readonly attribute ClassUList all_of_type_class;
    readonly attribute ClassUList all_of_class_class;
    const string CLASS_CONTAINMENT_RULES =
        "org.omg.mof:constraint.model.class.class_containment_rules";
    const string ABSTRACT_CLASSES_CANNOT_BE_SINGLETON =
        "org.omg.mof:constraint.model.class.abstract_classes_cannot_be_singleton";

    Class create_class (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from GeneralizableElement */   in boolean is_root,
        /* from GeneralizableElement */   in boolean is_leaf,
        /* from GeneralizableElement */   in boolean is_abstract,
        /* from GeneralizableElement */   in ::Model::VisibilityKind visibility,
        /* from Class */                   in boolean is_singleton)
        raises (Reflective::MofError);
}; // end of interface ClassClass

```

```

interface Class : ClassClass, Classifier {
    boolean is_singleton ()
        raises (Reflective::MofError);
    void set_is_singleton (in boolean new_value)
        raises (Reflective::MofError);
}; // end of interface Class

typedef ::CORBA::TypeCode TypeDescriptor;

interface DataTypeClass : ClassifierClass {
    readonly attribute DataTypeUList all_of_type_data_type;
    readonly attribute DataTypeUList all_of_class_data_type;
    const string DATA_TYPE_CONTAINMENT_RULES =
        "org.omg.mof:constraint.model.data_type.data_type_containment_rules";
    const string THIS_TYPECODE_NOT_SUPPORTED =
        "org.omg.mof:constraint.model.data_type.this_typecode_not_supported";
    const string DATA_TYPES_HAVE_NO_SUPERTYPES =
        "org.omg.mof:constraint.model.data_type.data_types_have_no_supertypes";
    const string DATA_TYPES_CANNOT_BE_ABSTRACT =
        "org.omg.mof:constraint.model.data_type.data_types_cannot_be_abstract";

    DataType create_data_type (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from GeneralizableElement */   in boolean is_root,
        /* from GeneralizableElement */   in boolean is_leaf,
        /* from GeneralizableElement */   in boolean is_abstract,
        /* from GeneralizableElement */   in ::Model::VisibilityKind visibility,
        /* from DataType */               in TypeDescriptor type_code)
        raises (Reflective::MofError);
}; // end of interface DataTypeClass

interface DataType : DataTypeClass, Classifier {
    TypeDescriptor type_code ()
        raises (Reflective::MofError);
    void set_type_code (in TypeDescriptor new_value)
        raises (Reflective::MofError);
}; // end of interface DataType

interface TypeAliasClass : TypedElementClass {
    readonly attribute TypeAliasUList all_of_type_type_alias;
    readonly attribute TypeAliasUList all_of_class_type_alias;
    TypeAlias create_type_alias (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation)
        raises (Reflective::MofError);
}; // end of interface TypeAliasClass

interface TypeAlias : TypeAliasClass, TypedElement {
}; // end of interface TypeAlias

enum ScopeKind { instance_level, classifier_level };

interface FeatureClass : ModelElementClass {
    readonly attribute FeatureUList all_of_type_feature;

```

```

}; // end of interface FeatureClass

interface Feature : FeatureClass, ModelElement {
    ScopeKind scope ()
        raises (Reflective::MofError);
    void set_scope (in ScopeKind new_value)
        raises (Reflective::MofError);
    VisibilityKind visibility ()
        raises (Reflective::MofError);
    void set_visibility (in VisibilityKind new_value)
        raises (Reflective::MofError);
}; // end of interface Feature

const long UNBOUNDED = -1;
struct MultiplicityType {
    long lower;
    long upper;
    boolean is_ordered;
    boolean is_unique; };
const string LOWER_CANNOT_BE_NEGATIVE_OR_UNBOUNDED =
    "org.omg.mof:constraint.model.multiplicity_type.lower_cannot_be_negative_or_unbounded";
const string LOWER_CANNOT_EXCEED_UPPER =
    "org.omg.mof:constraint.model.multiplicity_type.lower_cannot_exceed_upper";
const string UPPER_MUST_BE_POSITIVE =
    "org.omg.mof:constraint.model.multiplicity_type.upper_must_be_positive";
const string MUST_BE_UNORDERED_NONUNIQUE =
    "org.omg.mof:constraint.model.multiplicity_type.must_be_unordered_nonunique";

interface StructuralFeatureClass : FeatureClass, TypedElementClass {
    readonly attribute StructuralFeatureUList all_of_type_structural_feature;
}; // end of interface StructuralFeatureClass

interface StructuralFeature : StructuralFeatureClass, Feature, TypedElement {
    MultiplicityType multiplicity ()
        raises (Reflective::MofError);
    void set_multiplicity (in MultiplicityType new_value)
        raises (Reflective::MofError);
    boolean is_changeable ()
        raises (Reflective::MofError);
    void set_is_changeable (in boolean new_value)
        raises (Reflective::MofError);
}; // end of interface StructuralFeature

interface MofAttributeClass : StructuralFeatureClass {
    readonly attribute MofAttributeUList all_of_type_mof_attribute;
    readonly attribute MofAttributeUList all_of_class_mof_attribute;

    MofAttribute create_mof_attribute (
        /* from ModelElement */      in ::Model::NameType name,
        /* from ModelElement */      in ::Model::AnnotationType annotation,
        /* from Feature */           in ::Model::ScopeKind scope,
        /* from Feature */           in ::Model::VisibilityKind visibility,
        /* from StructuralFeature */ in ::Model::MultiplicityType multiplicity,
        /* from StructuralFeature */ in boolean is_changeable,
        /* from MofAttribute */      in boolean is_derived)

```



```

        raises (Reflective::MofError);
}; // end of interface MofAttributeClass

interface MofAttribute : MofAttributeClass, StructuralFeature {
    boolean is_derived ()
        raises (Reflective::MofError);
    void set_is_derived (in boolean new_value)
        raises (Reflective::MofError);
}; // end of interface MofAttribute

interface ReferenceClass : StructuralFeatureClass {
    readonly attribute ReferenceUList all_of_type_reference;
    readonly attribute ReferenceUList all_of_class_reference;
    const string REFERENCE_MULTIPLICITY_MUST_MATCH_END =
        "org.omg.mof:constraint.model.reference.reference_multiplicity_must_match_end";
    const string REFERENCE_MUST_BE_INSTANCE_SCOPED =
        "org.omg.mof:constraint.model.reference.reference_must_be_instance_scoped";
    const string CHANGEABLE_REFERENCE_MUST_HAVE_CHANGEABLE_END =
        "org.omg.mof:constraint.model.reference.changeable_reference_must_have_changeable_end";
    const string REFERENCE_TYPE_MUST_MATCH_END_TYPE =
        "org.omg.mof:constraint.model.reference.reference_type_must_match_end_type";
    const string REFERENCED_END_MUST_BE_NAVIGABLE =
        "org.omg.mof:constraint.model.reference.referenced_end_must_be_navigable";
    const string CONTAINER_MUST_MATCH_EXPOSED_TYPE =
        "org.omg.mof:constraint.model.reference.container_must_match_exposed_type";
    const string REFERENCED_END_MUST_BE_VISIBLE =
        "org.omg.mof:constraint.model.reference.referenced_end_must_be_visible";

    Reference create_reference (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Feature */                 in ::Model::ScopeKind scope,
        /* from Feature */                 in ::Model::VisibilityKind visibility,
        /* from StructuralFeature */       in ::Model::MultiplicityType multiplicity,
        /* from StructuralFeature */       in boolean is_changeable)
        raises (Reflective::MofError);
}; // end of interface ReferenceClass

interface Reference : ReferenceClass, StructuralFeature {
    AssociationEnd exposed_end ()
        raises (Reflective::MofError);
    void set_exposed_end (in AssociationEnd new_value)
        raises (Reflective::MofError);
    AssociationEnd referenced_end ()
        raises (Reflective::MofError);
    void set_referenced_end (in AssociationEnd new_value)
        raises (Reflective::MofError);
}; // end of interface Reference

interface BehavioralFeatureClass : FeatureClass, NamespaceClass {
    readonly attribute BehavioralFeatureUList all_of_type_behavioral_feature;
}; // end of interface BehavioralFeatureClass

interface BehavioralFeature :
    BehavioralFeatureClass, Feature , Namespace {};

```

```

interface OperationClass : BehavioralFeatureClass {
  readonly attribute OperationUList all_of_type_operation;
  readonly attribute OperationUList all_of_class_operation;

  const string OPERATION_CONTAINMENT_RULES =
    "org.omg.mof:constraint.model.operation.operation_containment_rules";
  const string OPERATIONS_HAVE_AT_MOST_ONE_RETURN =
    "org.omg.mof:constraint.model.operation.operations_have_at_most_one_return";
  const string OPERATION_EXCEPTIONS_MUST_BE_VISIBLE =
    "org.omg.mof:constraint.model.operation.operation_exceptions_must_be_visible";

  Operation create_operation (
    /* from ModelElement */           in ::Model::NameType name,
    /* from ModelElement */           in ::Model::AnnotationType annotation,
    /* from Feature */                 in ::Model::ScopeKind scope,
    /* from Feature */                 in ::Model::VisibilityKind visibility,
    /* from Operation */               in boolean is_query)
    raises (Reflective::MofError);
}; // end of interface OperationClass

interface Operation : OperationClass, BehavioralFeature {
  boolean is_query ()
    raises (Reflective::MofError);
  void set_is_query (in boolean new_value)
    raises (Reflective::MofError);
  MofExceptionUList exceptions ()
    raises (Reflective::MofError);
  void set_exceptions (in MofExceptionUList new_value)
    raises (Reflective::MofError);
  void add_exceptions (in MofException new_element)
    raises (Reflective::MofError);
  void add_exceptions_before (in MofException new_element,
                             in MofException before_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_exceptions (in MofException old_element,
                         in MofException new_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove_exceptions (in MofException old_element)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Operation

interface MofExceptionClass : BehavioralFeatureClass {
  readonly attribute MofExceptionUList all_of_type_mof_exception;
  readonly attribute MofExceptionUList all_of_class_mof_exception;
  const string EXCEPTION_CONTAINMENT_RULES =
    "org.omg.mof:constraint.model.mof_exception.exception_containment_rules";
  const string EXCEPTIONS_HAVE_ONLY_OUT_PARAMETERS =
    "org.omg.mof:constraint.model.mof_exception.exceptions_have_only_out_parameters";

  MofException create_mof_exception (
    /* from ModelElement */           in ::Model::NameType name,
    /* from ModelElement */           in ::Model::AnnotationType annotation,
    /* from Feature */                 in ::Model::ScopeKind scope,
    /* from Feature */                 in ::Model::VisibilityKind visibility)

```

```

        raises (Reflective::MofError);
}; // end of interface MofExceptionClass

interface MofException : MofExceptionClass , BehavioralFeature { };

interface AssociationClass : ClassifierClass {
    readonly attribute AssociationUList all_of_type_association;
    readonly attribute AssociationUList all_of_class_association;
    const string ASSOCIATIONS_CONTAINMENT_RULES =
        "org.omg.mof:constraint.model.association.associations_containment_rules";
    const string ASSOCIATIONS_HAVE_NO_SUPERTYPES =
        "org.omg.mof:constraint.model.association.associations_have_no_supertypes";
    const string ASSOCIATIONS_MUST_BE_ROOT_AND_LEAF =
        "org.omg.mof:constraint.model.association.associations_must_be_root_and_leaf";
    const string ASSOCIATIONS_CANNOT_BE_ABSTRACT =
        "org.omg.mof:constraint.model.association.associations_cannot_be_abstract";
    const string ASSOCIATIONS_MUST_BE_PUBLIC =
        "org.omg.mof:constraint.model.association.associations_must_be_public";
    const string ASSOCIATIONS_MUST_BE_BINARY =
        "org.omg.mof:constraint.model.association.associations_must_be_binary";

    Association create_association (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from GeneralizableElement */   in boolean is_root,
        /* from GeneralizableElement */   in boolean is_leaf,
        /* from GeneralizableElement */   in boolean is_abstract,
        /* from GeneralizableElement */   in ::Model::VisibilityKind visibility,
        /* from Association */             in boolean is_derived)
        raises (Reflective::MofError);
}; // end of interface AssociationClass

interface Association : AssociationClass, Classifier {

    boolean is_derived ()
        raises (Reflective::MofError);
    void set_is_derived (in boolean new_value)
        raises (Reflective::MofError);
}; // end of interface Association

enum AggregationKind { none, shared, composite };

interface AssociationEndClass : TypedElementClass {
    readonly attribute AssociationEndUList all_of_type_association_end;
    readonly attribute AssociationEndUList all_of_class_association_end;
    const string END_TYPE_MUST_BE_CLASS =
        "org.omg.mof:constraint.model.association_end.end_type_must_be_class";
    const string ENDS_MUST_BE_UNIQUE =
        "org.omg.mof:constraint.model.association_end.ends_must_be_unique";
    const string CANNOT_HAVE_TWO_ORDERED_ENDS =
        "org.omg.mof:constraint.model.association_end.cannot_have_two_ordered_ends";
    const string CANNOT_HAVE_TWO_AGGREGATE_ENDS =
        "org.omg.mof:constraint.model.association_end.cannot_have_two_aggregate_ends";

    AssociationEnd create_association_end (

```

```

        /* from ModelElement */
        /* from ModelElement */
        /* from AssociationEnd */
        /* from AssociationEnd */
        /* from AssociationEnd */
        /* from AssociationEnd */
        raises (Reflective::MofError);
}; // end of interface AssociationEndClass

interface AssociationEnd : AssociationEndClass, TypedElement {
    boolean is_navigable ()
        raises (Reflective::MofError);
    void set_is_navigable (in boolean new_value)
        raises (Reflective::MofError);
    AggregationKind aggregation ()
        raises (Reflective::MofError);
    void set_aggregation (in AggregationKind new_value)
        raises (Reflective::MofError);
    MultiplicityType multiplicity ()
        raises (Reflective::MofError);
    void set_multiplicity (in MultiplicityType new_value)
        raises (Reflective::MofError);
    AssociationEnd other_end ()
        raises (Reflective::MofError);
    boolean is_changeable ()
        raises (Reflective::MofError);
    void set_is_changeable (in boolean new_value);
}; // end of interface AssociationEnd

interface PackageClass : GeneralizableElementClass {
    readonly attribute PackageUList all_of_type_package;
    readonly attribute PackageUList all_of_class_package;
    const string PACKAGE_CONTAINMENT_RULES =
        "org.omg.mof.constraint.model.package.package_containment_rules";
    const string PACKAGES_CANNOT_BE_ABSTRACT =
        "org.omg.mof.constraint.model.package.packages_cannot_be_abstract";

    typedef string FormatType;

    exception FormatNotSupported {};
    exception ObjectNotExternalizable { string explanation; };
    exception IllformedExternalizedObject { string explanation; };

    GeneralizableElement internalize (in PackageClass::FormatType format,
        in any stream)
        raises (PackageClass::FormatNotSupported,
            PackageClass::IllformedExternalizedObject,
            Reflective::MofError);

    Package create_package (
        /* from ModelElement */
        /* from ModelElement */
        /* from GeneralizableElement */
        /* from GeneralizableElement */
        /* from GeneralizableElement */
        in ::Model::NameType name,
        in ::Model::AnnotationType annotation,
        in boolean is_root,
        in boolean is_leaf,
        in boolean is_abstract,

```

```

        /* from GeneralizableElement */           in ::Model::VisibilityKind visibility)
        raises (Reflective::MofError);
}; // end of interface PackageClass

interface Package : PackageClass, GeneralizableElement {

    any externalize (in PackageClass::FormatType format)
        raises (PackageClass::ObjectNotExternalizable,
                PackageClass::FormatNotSupported,
                Reflective::MofError);
}; // end of interface Package

interface ImportClass : ModelElementClass {
    readonly attribute ImportUList all_of_type_import;
    readonly attribute ImportUList all_of_class_import;
    const string IMPORTED_NAMESPACE_MUST_BE_VISIBLE =
        "org.omg.mof:constraint.model.imported_namespace_must_be_visible";
    const string CAN_ONLY_IMPORT_PACKAGES_AND_CLASSES =
        "org.omg.mof:constraint.model.import.can_only_import_packages_and_classes";
    const string CANNOT_IMPORT_SELF =
        "org.omg.mof:constraint.model.import.cannot_import_self";

    const string CANNOT_IMPORT_NESTED_COMPONENTS =
        "org.omg.mof:constraint.model.import.cannot_import_nested_components";
    const string NESTED_PACKAGES_CANNOT_IMPORT =
        "org.omg.mof:constraint.model.import.nested_packages_cannot_import";
    Import create_import (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Import */                  in ::Model::VisibilityKind visibility,
        /* from Import */                  in boolean is_clustered)
        raises (Reflective::MofError);
}; // end of interface ImportClass

interface Import : ImportClass, ModelElement {
    VisibilityKind visibility ()
        raises (Reflective::MofError);
    void set_visibility (in VisibilityKind new_value)
        raises (Reflective::MofError);
    boolean is_clustered ()
        raises (Reflective::MofError);
    void set_is_clustered (in boolean new_value)
        raises (Reflective::MofError);
    Namespace imported_namespace ()
        raises (Reflective::MofError);
    void set_imported_namespace (in Namespace new_value)
        raises (Reflective::MofError);
}; // end of interface Import

enum DirectionKind { in_dir, out_dir, inout_dir, return_dir };
interface ParameterClass : TypedElementClass {
    readonly attribute ParameterUList all_of_type_parameter;
    readonly attribute ParameterUList all_of_class_parameter;
    Parameter create_parameter (
        /* from ModelElement */           in ::Model::NameType name,

```

```

        /* from ModelElement */      in ::Model::AnnotationType annotation,
        /* from Parameter */         in ::Model::DirectionKind direction,
        /* from Parameter */         in ::Model::MultiplicityType multiplicity)
    raises (Reflective::MofError);
}; // end of interface ParameterClass

interface Parameter : ParameterClass, TypedElement {
    DirectionKind direction ()
        raises (Reflective::MofError);
    void set_direction (in DirectionKind new_value)
        raises (Reflective::MofError);
    MultiplicityType multiplicity ()
        raises (Reflective::MofError);
    void set_multiplicity (in MultiplicityType new_value)
        raises (Reflective::MofError);
}; // end of interface Parameter

interface ConstraintClass : ModelElementClass {
    readonly attribute ConstraintUList all_of_type_constraint;
    readonly attribute ConstraintUList all_of_class_constraint;
    const string CANNOT_CONSTRAIN_THIS_ELEMENT =
        "org.omg.mof.constraint.model.constraint.cannot_constrain_this_element";
    const string CONSTRAINTS_LIMITED_TO_CONTAINER =
        "org.omg.mof.constraint.model.constraint.constraints_limited_to_container";

    enum EvaluationKind { immediate, deferred };

    Constraint create_constraint (
        /* from ModelElement */      in ::Model::NameType name,
        /* from ModelElement */      in ::Model::AnnotationType annotation,
        /* from Constraint */         in any expression,
        /* from Constraint */         in string language,
        /* from Constraint */         in ::Model::ConstraintClass::EvaluationKind evaluation_policy)
        raises (Reflective::MofError);
}; // end of interface ConstraintClass

interface Constraint : ConstraintClass, ModelElement {
    any expression ()
        raises (Reflective::MofError);
    void set_expression (in any new_value)
        raises (Reflective::MofError);
    string language ()
        raises (Reflective::MofError);
    void set_language (in string new_value)
        raises (Reflective::MofError);
    ConstraintClass::EvaluationKind evaluation_policy ()
        raises (Reflective::MofError);
    void set_evaluation_policy (in ConstraintClass::EvaluationKind new_value)
        raises (Reflective::MofError);
    ModelElementSet constrained_elements ()
        raises (Reflective::MofError);
    void set_constrained_elements (in ModelElementSet new_value)
        raises (Reflective::MofError);
    void add_constrained_elements (in ModelElement new_element)

```

```

    raises (Reflective::MofError);
void modify_constrained_elements (in ModelElement old_element,
                                  in ModelElement new_element)
    raises (Reflective::MofError);
void remove_constrained_elements (in ModelElement old_element)
    raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Constraint

typedef any LiteralType;

interface ConstantClass : TypedElementClass {
    readonly attribute ConstantUList all_of_type_constant;
    readonly attribute ConstantUList all_of_class_constant;
    const string CONSTANTS_VALUE_MUST_MATCH_TYPE =
        "org.omg.mof:constraint.model.constant.constants_value_must_match_type";
    const string CONSTANTS_TYPE_MUST_BE_SIMPLE_DATA_TYPE =
        "org.omg.mof:constraint.model.constant.constants_type_must_be_simple_data_type";

    Constant create_constant (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Constant */                in ::Model::LiteralType const_value)
        raises (Reflective::MofError);
}; // end of interface ConstantClass

interface Constant : ConstantClass, TypedElement {
    LiteralType value ()
        raises (Reflective::MofError);
    void set_value (in LiteralType new_value)
        raises (Reflective::MofError);
}; // end of interface Constant

interface TagClass : ModelElementClass {
    readonly attribute TagUList all_of_type_tag;
    readonly attribute TagUList all_of_class_tag;

    Tag create_tag (
        /* from ModelElement */           in ::Model::NameType name,
        /* from ModelElement */           in ::Model::AnnotationType annotation,
        /* from Tag */                    in string tag_id,
        /* from Tag */                    in AnyBag values)
        raises (Reflective::MofError);
}; // end of interface TagClass

interface Tag : TagClass, ModelElement {
    string tag_id ()
        raises (Reflective::MofError);
    void set_tag_id (in string new_value)
        raises (Reflective::MofError);
    AnyBag values ()
        raises (Reflective::MofError);
    void set_values (in AnyBag new_value)
        raises (Reflective::MofError);
    void add_values (in any new_element)
        raises (Reflective::MofError);
};

```

```

void modify_values (in any old_element, in any new_element)
  raises (Reflective::NotFound, Reflective::MofError);
void remove_values (in any old_element)
  raises (Reflective::NotFound, Reflective::MofError);
ModelElementSet elements ()
  raises (Reflective::MofError);
void set_elements (in ModelElementSet new_value)
  raises (Reflective::MofError);
void add_elements (in ModelElement new_element)
  raises (Reflective::MofError);
void add_elements_before (in ModelElement new_element,
                          in ModelElement before_element)
  raises (Reflective::NotFound, Reflective::MofError);
void modify_elements (in ModelElement old_element,
                      in ModelElement new_element)
  raises (Reflective::NotFound, Reflective::MofError);
void remove_elements (in ModelElement old_element)
  raises (Reflective::NotFound, Reflective::MofError);
}; // end of interface Tag

struct AttachesToLink {
  ModelElement model_element;
  ::Model::Tag tag;
};
typedef sequence <AttachesToLink> AttachesToLinkSet;

interface AttachesTo : Reflective::RefAssociation {
  AttachesToLinkSet all_attaches_to_links ();
  boolean exists (in ModelElement model_element, in ::Model::Tag tag)
    raises (Reflective::MofError);
  ModelElementSet model_element (in ::Model::Tag tag)
    raises (Reflective::MofError);
  TagUList tag (in ModelElement model_element)
    raises (Reflective::MofError);
  void add (in ModelElement model_element, in ::Model::Tag tag)
    raises (Reflective::MofError);
  void add_before_tag (in ModelElement model_element,
                      in ::Model::Tag tag,
                      in ::Model::Tag before)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_model_element (in ModelElement model_element,
                             in ::Model::Tag tag,
                             in ModelElement new_model_element)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_tag (in ModelElement model_element,
                  in ::Model::Tag tag,
                  in ::Model::Tag new_tag)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (in ModelElement model_element, in ::Model::Tag tag)
    raises (Reflective::NotFound, Reflective::MofError);
};
struct DependsOnLink {
  ModelElement dependent;
  ModelElement provider;
};

```



```

typedef sequence <DependsOnLink> DependsOnLinkSet;

interface DependsOn : Reflective::RefAssociation {
    DependsOnLinkSet all_depends_on_links ();
    boolean exists (in ModelElement dependent, in ModelElement provider)
        raises (Reflective::MofError);
    ModelElementSet dependent (in ModelElement provider)
        raises (Reflective::MofError);
    ModelElementSet provider (in ModelElement dependent)
        raises (Reflective::MofError);
};

struct ContainsLink {
    Namespace container;
    ModelElement contained_element;
};
typedef sequence <ContainsLink> ContainsLinkSet;

interface Contains : Reflective::RefAssociation {
    ContainsLinkSet all_contains_links ();
    boolean exists (in Namespace container, in ModelElement contained_element)
        raises (Reflective::MofError);
    Namespace container (in ModelElement contained_element)
        raises (Reflective::MofError);
    ModelElementUList contained_element (in Namespace container)
        raises (Reflective::MofError);
    void add (in Namespace container, in ModelElement contained_element)
        raises (Reflective::MofError);
    void add_before_contained_element (in Namespace container,
                                       in ModelElement contained_element,
                                       in ModelElement before)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_container (in Namespace container,
                          in ModelElement contained_element,
                          in Namespace new_container)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_contained_element (in Namespace container,
                                   in ModelElement contained_element,
                                   in ModelElement new_contained_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (in Namespace container, in ModelElement contained_element)
        raises (Reflective::NotFound, Reflective::MofError);
};

struct GeneralizesLink {
    GeneralizableElement supertype;
    GeneralizableElement subtype;
};
typedef sequence <GeneralizesLink> GeneralizesLinkSet;

interface Generalizes : Reflective::RefAssociation {
    GeneralizesLinkSet all_generalizes_links ();
    boolean exists (in GeneralizableElement supertype,
                  in GeneralizableElement subtype)
        raises (Reflective::MofError);
};

```

```
GeneralizableElementUList supertype (in GeneralizableElement subtype)
  raises (Reflective::MofError);
GeneralizableElementSet subtype (in GeneralizableElement supertype)
  raises (Reflective::MofError);
void add (in GeneralizableElement supertype,
         in GeneralizableElement subtype)
  raises (Reflective::MofError);
void add_before_supertype (in GeneralizableElement supertype,
                          in GeneralizableElement subtype,
                          in GeneralizableElement before)
  raises (Reflective::NotFound, Reflective::MofError);
void modify_supertype (in GeneralizableElement supertype,
                     in GeneralizableElement subtype,
                     in GeneralizableElement new_supertype)
  raises (Reflective::NotFound, Reflective::MofError);
void modify_subtype (in GeneralizableElement supertype,
                   in GeneralizableElement subtype,
                   in GeneralizableElement new_subtype)
  raises (Reflective::NotFound, Reflective::MofError);
void remove (in GeneralizableElement supertype,
            in GeneralizableElement subtype)
  raises (Reflective::NotFound, Reflective::MofError);
};

struct AliasesLink {
  Import importer;
  Namespace imported;
};
typedef sequence <AliasesLink> AliasesLinkSet;

interface Aliases : Reflective::RefAssociation {
  AliasesLinkSet all_aliases_links ();
  boolean exists (in Import importer, in Namespace imported)
    raises (Reflective::MofError);
  ImportSet importer (in Namespace imported)
    raises (Reflective::MofError);
  Namespace imported (in Import importer)
    raises (Reflective::MofError);
  void add (in Import importer, in Namespace imported)
    raises (Reflective::MofError);
  void modify_importer (in Import importer,
                      in Namespace imported,
                      in Import new_importer)
    raises (Reflective::NotFound, Reflective::MofError);
  void modify_imported (in Import importer,
                      in Namespace imported,
                      in Namespace new_imported)
    raises (Reflective::NotFound, Reflective::MofError);
  void remove (in Import importer, in Namespace imported)
    raises (Reflective::NotFound, Reflective::MofError);
};

struct ConstrainsLink {
  ::Model::Constraint constraint;
  ModelElement constrained_element;
```

```

};
typedef sequence <ConstrainsLink> ConstrainsLinkSet;

interface Constrains : Reflective::RefAssociation {
    ConstrainsLinkSet all_constrains_links ();
    boolean exists (in ::Model::Constraint constraint,
                   in ModelElement constrained_element)
        raises (Reflective::MofError);
    ConstraintSet constraint (in ModelElement constrained_element)
        raises (Reflective::MofError);
    ModelElementSet constrained_element (in ::Model::Constraint constraint)
        raises (Reflective::MofError);
    void add (in ::Model::Constraint constraint,
             in ModelElement constrained_element)
        raises (Reflective::MofError);
    void modify_constraint (in ::Model::Constraint constraint,
                           in ModelElement constrained_element,
                           in ::Model::Constraint new_constraint)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_constrained_element (in ::Model::Constraint constraint,
                                     in ModelElement constrained_element,
                                     in ModelElement new_constrained_element)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (in ::Model::Constraint constraint,
                in ModelElement constrained_element)
        raises (Reflective::NotFound, Reflective::MofError);
};

struct CanRaiseLink {
    ::Model::Operation operation;
    MofException except;
};
typedef sequence <CanRaiseLink> CanRaiseLinkSet;

interface CanRaise : Reflective::RefAssociation {
    CanRaiseLinkSet all_can_raise_links ();
    boolean exists (in ::Model::Operation operation, in MofException except)
        raises (Reflective::MofError);
    OperationSet operation (in MofException except)
        raises (Reflective::MofError);
    MofExceptionUList except (in ::Model::Operation operation)
        raises (Reflective::MofError);
    void add (in ::Model::Operation operation, in MofException except)
        raises (Reflective::MofError);
    void add_before_except (in ::Model::Operation operation,
                           in MofException except,
                           in MofException before)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_operation (in ::Model::Operation operation,
                          in MofException except,
                          in ::Model::Operation new_operation)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_except (in ::Model::Operation operation,
                      in MofException except,
                      in MofException new_except)
};

```

```
        raises (Reflective::NotFound, Reflective::MofError);
void remove (in ::Model::Operation operation, in MofException except)
        raises (Reflective::NotFound, Reflective::MofError);
};

struct ExposesLink {
    Reference referrer;
    AssociationEnd exposed_end;
};
typedef sequence <ExposesLink> ExposesLinkSet;

interface Exposes : Reflective::RefAssociation {
    ExposesLinkSet all_exposes_links ();
    boolean exists (in Reference referrer, in AssociationEnd exposed_end)
        raises (Reflective::MofError);
    ReferenceSet referrer (in AssociationEnd exposed_end)
        raises (Reflective::MofError);
    AssociationEnd exposed_end (in Reference referrer)
        raises (Reflective::MofError);
    void add (in Reference referrer, in AssociationEnd exposed_end)
        raises (Reflective::MofError);
    void modify_referrer (in Reference referrer,
                        in AssociationEnd exposed_end,
                        in Reference new_referrer)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_exposed_end (in Reference referrer,
                          in AssociationEnd exposed_end,
                          in AssociationEnd new_exposed_end)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (in Reference referrer, in AssociationEnd exposed_end)
        raises (Reflective::NotFound, Reflective::MofError);
};

struct RefersToLink {
    Reference referent;
    AssociationEnd referenced_end;
};
typedef sequence <RefersToLink> RefersToLinkSet;

interface RefersTo : Reflective::RefAssociation {
    RefersToLinkSet all_refers_to_links ();
    boolean exists (in Reference referent, in AssociationEnd referenced_end)
        raises (Reflective::MofError);
    ReferenceSet referent (in AssociationEnd referenced_end)
        raises (Reflective::MofError);
    AssociationEnd referenced_end (in Reference referent)
        raises (Reflective::MofError);
    void add (in Reference referent, in AssociationEnd referenced_end)
        raises (Reflective::MofError);
    void modify_referent (in Reference referent,
                       in AssociationEnd referenced_end,
                       in Reference new_referent)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_referenced_end (in Reference referent,
                              in AssociationEnd referenced_end,
```

```

        in AssociationEnd new_referenced_end)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (in Reference referent, in AssociationEnd referenced_end)
        raises (Reflective::NotFound, Reflective::MofError);
};

```

```

struct IsOfTypeLink {
    Classifier type;
    TypedElement typed_elements;
};
typedef sequence <IsOfTypeLink> IsOfTypeLinkSet;

```

```

interface IsOfType : Reflective::RefAssociation {
    IsOfTypeLinkSet all_is_of_type_links ();
    boolean exists (in Classifier type, in TypedElement typed_elements)
        raises (Reflective::MofError);
    Classifier type (in TypedElement typed_elements)
        raises (Reflective::MofError);
    TypedElementSet typed_elements (in Classifier type)
        raises (Reflective::MofError);
    void add (in Classifier type, in TypedElement typed_elements)
        raises (Reflective::MofError);
    void modify_type (in Classifier type,
                    in TypedElement typed_elements,
                    in Classifier new_type)
        raises (Reflective::NotFound, Reflective::MofError);
    void modify_typed_elements (in Classifier type,
                               in TypedElement typed_elements,
                               in TypedElement new_typed_elements)
        raises (Reflective::NotFound, Reflective::MofError);
    void remove (in Classifier type, in TypedElement typed_elements)
        raises (Reflective::NotFound, Reflective::MofError);
};

```

```

interface ModelPackageFactory {
    ModelPackage create_model_package ()
        raises (Reflective::MofError);
};

```

```

interface ModelPackage : Reflective::RefPackage {
    readonly attribute ModelElementClass model_element_ref;
    readonly attribute NamespaceClass namespace_ref;
    readonly attribute GeneralizableElementClass
        generalizable_element_ref;
    readonly attribute TypedElementClass typed_element_ref;
    readonly attribute ClassifierClass classifier_ref;
    readonly attribute ClassClass class_ref;
    readonly attribute DataTypeClass data_type_ref;
    readonly attribute TypeAliasClass type_alias_ref;
    readonly attribute FeatureClass feature_ref;
    readonly attribute StructuralFeatureClass structural_feature_ref;
    readonly attribute MofAttributeClass mof_attribute_ref;
    readonly attribute ReferenceClass reference_ref;
    readonly attribute BehavioralFeatureClass behavior_feature_ref;
};

```

```

readonly attribute OperationClass operation_ref;
readonly attribute MofExceptionClass mof_exception_ref;
readonly attribute AssociationClass association_ref;
readonly attribute AssociationEndClass association_end_ref;
readonly attribute PackageClass package_ref;
readonly attribute ImportClass import_ref;
readonly attribute ParameterClass parameter_ref;
readonly attribute ConstraintClass constraint_ref;
readonly attribute ConstantClass constant_ref;
readonly attribute TagClass tag_ref;
readonly attribute AttachesTo attaches_to_ref;
readonly attribute DependsOn depends_on_ref;
readonly attribute Contains contains_ref;
readonly attribute Generalizes generalizes_ref;
readonly attribute Aliases aliases_ref;
readonly attribute Constrains constrains_ref;
readonly attribute CanRaise can_raise_ref;
readonly attribute Exposes exposes_ref;
readonly attribute RefersTo refers_to_ref;
readonly attribute IsOfType is_of_type_ref;
};
}; // end of module Model

// end of IDL generation

```

B.2 Reflective IDL

```

#pragma prefix "org.omg.mof"
module Reflective {
    interface RefBaseObject;
    interface RefObject;
    typedef sequence < RefObject > RefObjectUList;
    typedef sequence < RefObject > RefObjectSet;
    interface RefAssociation;
    interface RefPackage;
    typedef RefObject DesignatorType;
    typedef any ValueType;
    typedef sequence < ValueType > ValueTypeList;
    typedef sequence < RefObject, 2 > Link;
    typedef sequence < Link > LinkSet;
    const string UNDERFLOW_VIOLATION = "org.omg.mof:structural.underflow";
    const string OVERFLOW_VIOLATION = "org.omg.mof:structural.overflow";
    const string DUPLICATE_VIOLATION = "org.omg.mof:structural.duplicate";
    const string REFERENCE_CLOSURE_VIOLATION = "org.omg.mof:structural.reference_closure";
    const string SUPERTYPE_CLOSURE_VIOLATION = "org.omg.mof:structural.supertype_closure";
    const string COMPOSITION_CYCLE_VIOLATION = "org.omg.mof:structural.composition_cycle";
    const string COMPOSITION_CLOSURE_VIOLATION = "org.omg.mof:structural.composition_closure";
    const string INVALID_OBJECT_VIOLATION = "org.omg.mof:structural.invalid_object";
    const string NIL_OBJECT_VIOLATION = "org.omg.mof:structural.nil_object";
    const string INACCESSIBLE_OBJECT_VIOLATION = "org.omg.mof:structural.inaccessible_object";
    const string ALREADY_EXISTS_VIOLATION = "org.omg.mof:structural.already_exists";
    const string INVALID_DESIGNATOR_VIOLATION = "org.omg.mof:reflective.invalid_designator";
    const string WRONG_DESIGNATOR_DESIGNATOR_VIOLATION =
        "org.omg.mof:reflective.wrong_designator_kind";
}

```

```

const string UNKNOWN_DESIGNATOR_VIOLATION = "org.omg.mof:reflective.unknown_designator";
const string ABSTRACT_CLASS_VIOLATION = "org.omg.mof:reflective.abstract_class";
const string NOT_CHANGEABLE_VIOLATION = "org.omg.mof:reflective.not_changeable";
const string NOT_PUBLIC_VIOLATION = "org.omg.mof:reflective.not_public";
const string WRONG_SCOPE_VIOLATION = "org.omg.mof:reflective.wrong_scope";
const string WRONG_MULTPLICITY_VIOLATION = "org.omg.mof:reflective.wrong_multiplicity";
const string WRONG_TYPE_VIOLATION = "org.omg.mof:reflective.wrong_type";
const string WRONG_NUMBER_PARAMETERS_VIOLATION =
    "org.omg.mof:reflective.wrong_number_parameters";
const string INVALID_DELETION_VIOLATION = "org.omg.mof:reflective.invalid_deletion";

```

```

struct NamedValueType {
    string name;
    ValueType value;
};
typedef sequence < NamedValueType > NamedValueList;
exception MofError {
    string error_kind;
    RefObject element_in_error;
    NamedValueList extra_info;
    string error_description;
};
exception NotFound {};
exception NotSet {};
exception BadPosition {
    unsigned long current_size;
};
exception OtherException {
    DesignatorType exception_designator;
    ValueTypeList exception_values;
};

```

```

interface RefBaseObject {
    string ref_mof_id ();
    DesignatorType ref_meta_object ();
    boolean ref_itself (in RefBaseObject other_object);
    RefPackage ref_immediate_package ();
    RefPackage ref_outermost_package ();
    void ref_delete ()
        raises (MofError);
}; // end of RefBaseObject interface

```

```

interface RefObject : RefBaseObject {
    boolean ref_is_instance_of (in DesignatorType some_class,
        in boolean consider_subtypes);
    RefObject ref_create_instance (in ValueTypeList args)
        raises (MofError);
    RefObjectSet ref_all_objects (in boolean include_subtypes);
    void ref_set_value (in DesignatorType feature,
        in ValueType new_value)
        raises (MofError);
    ValueType ref_value (in DesignatorType feature)
        raises (NotSet, MofError);
    void ref_unset_value ()
        raises (MofError);
};

```

```

void ref_add_value (in DesignatorType feature,
                  in ValueType new_element)
    raises (MofError);
void ref_add_value_before (in DesignatorType feature,
                          in ValueType new_element,
                          in ValueType before_element)
    raises (NotFound, MofError);
void ref_add_value_at (in DesignatorType feature,
                    in ValueType new_element,
                    in unsigned long position)
    raises (BadPosition, MofError);
void ref_modify_value (in DesignatorType feature,
                    in ValueType old_element,
                    in ValueType new_element)
    raises (NotFound, MofError);
void ref_modify_value_at (in DesignatorType feature,
                        in ValueType new_element,
                        in unsigned long position)
    raises (BadPosition, MofError);
void ref_remove_value (in DesignatorType feature,
                    in ValueType existing_element)
    raises (NotFound, MofError);
void ref_remove_value_at (in DesignatorType feature,
                        in unsigned long position)
    raises (BadPosition, MofError);
RefObject ref_immediate_composite ();
RefObject ref_outermost_composite ();
ValueType ref_invoke_operation (in DesignatorType requested_operation,
                              inout ValueTypeList args)
    raises (OtherException, MofError);
}; // end of RefObject interface

```

```

interface RefAssociation : RefBaseObject {
    LinkSet ref_all_links ();
    boolean ref_link_exists (in Link some_link)
        raises (MofError);
    RefObjectUList ref_query (in DesignatorType query_end,
                            in RefObject query_object)
        raises (MofError);
    void ref_add_link (in Link new_link)
        raises (MofError);
    void ref_add_link_before (in Link new_link,
                            in DesignatorType position_end,
                            in RefObject before)
        raises (NotFound, MofError);
    void ref_modify_link (in Link old_link,
                        in DesignatorType position_end,
                        in RefObject new_object)
        raises (NotFound, MofError);
    void ref_remove_link (in Link old_link)
        raises (NotFound, MofError);
}; // end of RefAssociation interface

```

```

interface RefPackage : RefBaseObject {
    RefObject ref_class_ref (in DesignatorType class)

```



```
    raises (MofError);
RefAssociation ref_association_ref (in DesignatorType association)
    raises (MofError);
RefPackage ref_package_ref (in DesignatorType package)
    raises (MofError);
}; // end of RefPackage interface
}; // end of Reflective module
```


MODL Description of the MOF

C

The CORBA IDL for the Model module listed in Appendix B was automatically produced using a prototype MOF and associated tools developed by DSTC. This appendix gives the input files used to drive the IDL generation process. These are expressed in an interim version of the DSTC's Meta Object Definition language abbreviated as MODL.

MODL provides users with a compilable textual language to express models using the concepts of the MOF Model. It has a syntax based loosely on CORBA IDL, that has a direct correspondence with MOF Model concepts. For historical reasons, some of the MODL constructs use different names. In particular, References are called "knowns."

A rough specification of the interim MODL language is available on the OMG FTP server and DSTC's Meta Object Facility web page at:

<http://www.dstc.edu.au/Meta-Object-Facility/Review/>

C.1 MOF Model

```
//  
// MOF model expressed in MODL-2  
//  
#pragma idl_prefix "org.omg.mof"  
package Model {  
// This artifact type represents an element name. It should  
// conform to the CORBA identifier syntax rules, or start with  
// the character "*". The latter case should only be used for  
// anonymous data types (e.g. 'unsigned long') and for an  
// operation's return parameter.  
typedef string NameType;
```

```
// This artifact type represents annotations that are used
// to document the nodes of a meta-model.
typedef string AnnotationType;

// All elements of the core MOF meta-meta-model are derived
// from ModelElement. It provides all elements with a name
// and an annotation, along with "containedness", "constraints"
// and derived "dependency.
abstract class ModelElement {
    constraint MustBeContainedUnlessPackage : "[C-1]";
    constraint FrozenAttributesCannotBeChanged : "[C-2]";
    constraint FrozenElementsCannotBeDeleted : "[C-3]";
    constraint FrozenDependenciesCannotBeChanged : "[C-4]";

    // ***** MODEL ELEMENT NAMES *****

    // The model element's simple name
    attribute NameType name;
// The model element's fully qualified name

    readonly derived attribute list [1..*] of NameType qualified_name;

// ***** ANNOTATIONS *****

    // The model element's annotation is a multi-line string
    attribute AnnotationType annotation;

    // ***** DEPENDENCIES *****

    // Each model element is aware of what it depends on.
    known required_elements projects dependent of DependsOn;

    // This artifact type describes the kinds of dependency.
    typedef string DependencyKind;

    const DependencyKind ContainerDep = "container";
    const DependencyKind ContentsDep = "contents";
    const DependencyKind SignatureDep = "signature";
    const DependencyKind ConstraintDep = "constraint";
    const DependencyKind ConstrainedElementsDep = "constrained elements";
    const DependencyKind SpecializationDep = "specialization";
```

```
const DependencyKind ImportDep = "import";
const DependencyKind TypeDefinitionDep = "type definition";
const DependencyKind ReferencedEndsDep = "referenced ends";
const DependencyKind TaggedElementsDep = "tagged elements";
const DependencyKind IndirectDep = "indirect";
const DependencyKind AllDep = "all";

// Return this ME's set of dependents of the requested kind.
set [0..*] of ModelElement
  find_required_elements(
    // list of dependency kinds of interest
    in set [1..*] of DependencyKind kinds,
    // If true, return the closure of the dependency
    in boolean recursive);

// This operation categorizes the dependency of this ME on another.
// The result is false (and kind is "") if there is no dependency.
boolean is_required_because(
  // the model element (supposedly) depended on
  in ModelElement other_element,
  // the kind of dependency
  out DependencyKind reason);

// **** CONTAINMENT ****

// Each model element is aware of its container
known container projects contained_element of Contains;

// **** CONSTRAINTS, VERIFICATION and FREEZING ****

known constraints projects constrained_element of Constrains;

// A VerifyResult says whether or not a model element is valid, and
// if valid whether the caller can rely on this continuing to be the
// case.
enum VerifyResultKind {
  valid, published, invalid
};

// The Depth argument selects shallow or deep verification.
enum DepthKind {
  shallow, deep
};
```

```
// The ViolationType struct is used to return information about
// a constraint violation detected during the verify operation.
// The contents of this struct are (deliberately) identical to the contents of
// the Reflective::MofError exception.
struct ViolationType {
    string error_kind;
    Reflective::RefObject element_in_error;
    Reflective::NamedValueList values_in_error;
    string error_description;
};

// This operation verifies the model-element and its contents
// in context. The result will be true if everything is
// semantically consistent. Otherwise, error reports will
// be returned via "problems".
VerifyResultKind verify(
    // verify just this object or the closure
    // of the objects it depends on
    in DepthKind depth,
    // any errors found are returned via this param
    out set [0..*] of ViolationType problems);

// Returns true iff this model-element has been frozen
boolean is_frozen();

// **** VISIBILITY ****

// This operation checks whether the supplied ModelElement is visible
// to this one according to the visibility rules.
boolean is_visible(in ModelElement other_element);

};

// This artifact type represents the visibility of a component
// or feature beyond its defining context (i.e. its enclosing
// containers). Private means there is no visibility, protected
// means that there is visibility via generalization, and public
// means that there is visibility via importing as well.
enum VisibilityKind {public_vis, protected_vis, private_vis};

//
```

```

// Forward declaration of classes used in Namespace.
//
abstract class Namespace : ModelElement;
abstract class GeneralizableElement : Namespace;
abstract class Classifier : GeneralizableElement;
class quote "Class" : Classifier;

// Namespace represents any meta-meta-model element that
// acts as a container for other elements; e.g. has components
// or features. It has an associated derived namespace for
// the contained elements.
//
// An Namespace can also be imported into another Package, though the
// semantics are unclear, and the mapping to CORBA IDL is problematic.
abstract class Namespace : ModelElement {
    constraint ContentNamesMustNotCollide : "[C-5]";

    exception NameNotFound {
        NameType name;
    };

    exception NameNotResolved {
        string explanation;
        list [0..*] of NameType restOfName;
    };

known contents projects container of Contains;

// This op looks for an entry called 'name' in 'contents'
ModelElement lookup_element(in NameType name)
    raises (NameNotFound);

// This op does a name resolution of 'compound' in 'contents'
ModelElement
    resolve_qualified_name(in list [1..*] of NameType qualified_name)
    raises (NameNotResolved);
// This op returns the Model Elements in the Namespace's direct
// contents that match "of_type". If "include_subtypes" is true, the
// result includes ME's that are instances of subtypes of "of_type"
ordered set [0..*] of ModelElement
    find_elements_by_type(in Class of_type,
        in boolean include_subtypes);

```

```
// This op determines whether the supplied name would be valid
// name for a new ModelElement in the NameSpace.
boolean name_is_valid(in NameType proposed_name);

};

// GeneralizableElement represents a meta-meta-model element
// that can inherit from another one of the same kind, via the
// Generalizes association.
abstract class GeneralizableElement : Namespace {
    constraint SupertypeMustNotBeSelf : "[C-6]";
    constraint SupertypeKindMustBeSame : "[C-7]";
    constraint ContentsMustNotCollideWithSupertypes : "[C-8]";
    constraint DiamondRuleMustBeObeyed : "[C-9]";
    constraint NoSupertypesAllowedForRoot : "[C-10]";
    constraint SupertypesMustBeVisible : "[C-11]";
    constraint NoSubtypesAllowedForLeaf : "[C-12]";

    // If true, this GE cannot have supertypes
    attribute boolean is_root;

    // If true, this GE cannot have subtypes
    attribute boolean is_leaf;

// If true, this GE is abstract. This means that there will
// be no factory operation for the object.
    attribute boolean is_abstract;

    // This controls what can see the GE's in a namespace
    attribute VisibilityKind visibility;

    // A GE knows about its supertypes (but not its subtypes)
    known supertypes projects subtype of Generalizes;

    // This operation provides all of the GE's supertypes, not
    // just the immediate ones.
    ordered set [0..*] of GeneralizableElement all_supertypes();

    // Analogue to Namespace.lookup_element that includes the
    // contents of the GE's direct and indirect supertypes.
    ModelElement lookup_element_extended(in NameType name)
        raises (NameNotFound);
```



```
// Analogue to Namespace.find_elements_by_type that includes
// the contents of the GE's direct and indirect supertypes.
ordered set [0..*] of ModelElement
    find_elements_by_type_extended(in Class of_type,
                                  in boolean include_subtypes);
};

// TypedElement is an abstract subtype for those ModelElements
// that require a type as part of their definition.
abstract class TypedElement : ModelElement {
    constraint AssociationsCannotBeTypes : "[C-13]";
    constraint TypeMustBeVisible : "[C-14]";

    known type projects typed_elements of IsOfType;
};

// Classifier is the abstract superclass for things that can be
// the type of something else.
abstract class Classifier : GeneralizableElement {
};

// A Class represents the type of an object; i.e. a value with
// object identity. The class's interface is expressed using Operation,
// Attribute and AssociationEnd features.
class quote "Class" : Classifier {
    constraint ClassContainmentRules : "[C-15]";
    constraint AbstractClassesCannotBeSingleton : "[C-16]";

    // If the "is_singleton" attribute is true, the generated interfaces
    // will only allow a single instance of the class to exist.
    attribute boolean is_singleton;
};

typedef TypeCode TypeDescriptor;

// A DataType represents the type of a value that does not have
// object identity. It is used for expressing "artifact types"
// in the model level.
//
```

```
// DataTypes do not have any features, they may not generalize
// or be generalised, and they cannot be abstract.
Class DataType : Classifier {
    constraint DataTypeContainmentRules : "[C-17]";
    constraint ThisTypecodeNotSupported : "[C-18]";
    constraint DataTypesHaveNoSupertypes : "[C-19]";
    constraint DataTypesCannotBeAbstract : "[C-20]";

    // The described type can be any CORBA type and is expressed as
    // a CORBA typecode value.
    attribute TypeDescriptor type_code;
};

// A TypeAlias is a component of a DataType that is used to associate
// an embedded typecode within the DataType's typecode with another
// Classifier object.
class TypeAlias : TypedElement {
};

enum ScopeKind { instance_level, classifier_level };

// Feature is an abstract supertype for a number of kinds
// of "features" of classes and associations.
abstract class Feature : ModelElement {
    // Features may be "instance" or "class" level. Note however
    // that "class" level features are only allowed for features
    // of a Class.
    attribute ScopeKind scope;

    // A Feature may be "public", "private" or "protected". Note
    // that "private" and "protected" features are ignored for
    // IDL generation purposes.
    attribute VisibilityKind visibility;
};

const long Unbounded = -1;

// This artifact type is used to describe the number of values
// allowed / stored in a given context, and how they are related.
// The precise interpretation of this info depends on the context.
struct MultiplicityType {
    long lower;
```

```
    long upper;
    boolean is_ordered;
    boolean is_unique;
};
constraint LowerCannotBeNegativeOrUnbounded
    on MultiplicityType in MultiplicityType: "[C-54]";
constraint LowerCannotExceedUpper
    on MultiplicityType in MultiplicityType: "[C-55]";
constraint UpperMustBePositive
    on MultiplicityType in MultiplicityType: "[C-56]";
constraint MustBeUnorderedNonunique
    on MultiplicityType in MultiplicityType: "[C-57]";

// StructuralFeature is an abstract superclass for those
// features that are part of the "structure" of a Class; i.e.
// Attributes and References.
abstract class StructuralFeature : Feature, TypedElement {
    attribute MultiplicityType multiplicity;

// If is_changeable is true, the generated IDL will allow
// client updates of / through the attribute or association ref.
    attribute boolean is_changeable;
};

// An Attribute is a feature of a Class
#pragma idl_substitute_name "MofAttribute"
    class QUOTE "Attribute" : StructuralFeature {
attribute boolean is_derived;
    };

// A Reference is a feature of a Class that allows the client to
// treat a projection of some "known" Association involving this
// object as a navigable link. The Reference is linked to roles
// of an Association via the Exposes and RefersTo associations.
class Reference : StructuralFeature {
    constraint ReferenceMultiplicityMustMatchEnd : "[C-21]";
    constraint ReferenceMustBeInstanceScoped : "[C-22]";
    constraint ChangeableReferenceMustHaveChangeableEnd : "[C-23]";
    constraint ReferenceTypeMustMatchEndType : "[C-24]";
    constraint ReferencedEndMustBeNavigable : "[C-25]";
    constraint ContainerMustMatchExposedType : "[C-26]";
    constraint ReferencedEndMustBeVisible : "[C-27]";
```

```
known exposed_end projects referrer of Exposes;

known referenced_end projects referent of RefersTo;
};

// This class is the superclass of Operation and Exception
abstract class BehavioralFeature : Feature, Namespace {
};

// An Operation has Parameters as features, and is associated
// with the Exceptions that it raises.
class Operation : BehavioralFeature {
    constraint OperationContainmentRules : "[C-28]";
    constraint OperationsHaveAtMostOneReturn : "[C-29]";
    constraint OperationExceptionsMustBeVisible : "[C-30]";

    // If is_query is true, the operation should not alter the state of the
    // object to which the operation applies.
    attribute boolean is_query;

    // An operation knows about its exceptions
    known exceptions projects operation of CanRaise;
};

// An Exception has Parameters.
#pragma idl_substitute_name "MofException"
class QUOTE "Exception" : BehavioralFeature {

constraint ExceptionContainmentRules : "[C-31]";
    constraint ExceptionsHaveOnlyOutParameters : "[C-32]";
};

// An Association represents an relation between Classes
// The roles of the Association are described by AssociationEnd features.
// The MOF specification will say that only binary Associations need
// to be supported. Association attributes and operations are not
// supported in the core.
class QUOTE "Association" : Classifier {
    constraint AssociationContainmentRules : "[C-33]";
```

```
constraint AssociationsHaveNoSupertypes : "[C-34]";
constraint AssociationMustBeRootAndLeaf : "[C-35]";
constraint AssociationsCannotBeAbstract : "[C-36]";
constraint AssociationsMustBePublic : "[C-37]";
constraint AssociationsMustBeBinary : "[C-38]";

// This attribute says that the association information is
// derived from other information.
attribute boolean is_derived;
};

enum AggregationKind { none, QUOTE "shared", QUOTE "composite" };

// AssociationEnd is a feature of an Association that
// describes one of its roles; i.e. a column of the association
// table.
class AssociationEnd : TypedElement {
    constraint EndTypeMustBeClass : "[C-39]";
    constraint EndsMustBeUnique : "[C-40]";
    constraint EndsMustBeUnique : "[C-41]";
    constraint CannotHaveTwoAggregateEnds : "[C-42]";

    // Iff is_navigable is true, a Reference may Expose this role
    attribute boolean is_navigable;

    // The aggregation specifies the containment / sharing semantics of
    // the Role. (Does this belong on the Role or the Association?)
    attribute AggregationKind aggregation;

    // The multiplicity on a Role is a constraint on the number and
    // kind of values that fill the role when the association is
    // "projected" over a single value in the other role. [It is
    // not clear whether this is a meaningful definition given that
    // we may be allowing duplicate "rows" in the association "table".
    // Furthermore, it is not clear that either of is_ordered or
    // is_unique are well-defined under this definition.]
    attribute MultiplicityType multiplicity;

    // If is_changeable is true, the generated IDL will allow
    // client updates
    attribute boolean is_changeable;
};
```

```
// The "other end" of this Association
AssociationEnd other_end();
};

// A Package is a concrete model element that collects a number
// of related classes, associations, data types and constants.
// Packages may be nested, and may import objects from other
// Packages.
//
// There are a number of unresolved issues relating to the
// mapping of Package generalization and importing onto CORBA IDL.
class QUOTE "Package" : GeneralizableElement {
    constraint PackageContainmentRules : "[C-43]";
    constraint PackagesCannotBeAbstract : "[C-44]";

    // **** EXTERNALIZE / INTERNALIZE ****

    // A FormatType string denotes an externalization format
    typedef string FormatType;

    // This is raised if the caller requests an externalization
    // format that is not supported (or not known!) by this MOF
    // meta-meta-object implementation.
    exception FormatNotSupported {};

    // This is raised if the GE or its contents is in a state
    // that makes them unexternalizable; e.g. some externalization
    // formats may require that the GE is consistent.
    exception ObjectNotExternalizable {
        string explanation;
    };

    // This is raised if the externalized GE is ill-formed, or
    // of the wrong format.
    exception IllformedExternalizedObject {
        string explanation;
    };

    // The Externalize op converts the GE and its contents into
    // a "flat" form that can be passed by value. The 'format'
    // argument allows the client to select the externalization
    // format.
    any externalize(in FormatType format)
```

```
        raises (ObjectNotExternalizable, FormatNotSupported);

// The Internalize op creates a new GE from a "flat" form.
class Package internalize(in FormatType format, in any stream)
    raises (FormatNotSupported, IllformedExternalizedObject);
};

// A Import is a "feature" of a Package that refers
// to (imports) a component of another Package.
class QUOTE "Import" : ModelElement {
    constraint ImportedNamespaceMustBeVisible : "[C-45]";
    constraint CanOnlyImportPackagesAndClasses : "[C-46]";
    constraint CannotImportSelf : "[C-47]";
    constraint CannotImportNestedComponents : "[C-48]";
    constraint NestedPackagesCannotImport : "[C-49]";

    attribute VisibilityKind visibility;

    attribute boolean is_clustered;

        known imported_namespace projects importer of Aliases;
};

enum DirectionKind { in_dir, out_dir, inout_dir, return_dir };

// A Parameter is a component of an Operation or an Exception
// description.
class Parameter : TypedElement {
    // The parameters of an Operation may have direction "in",
    // "out" or "inout". The result of an Operation is expressed
    // as a Parameter with direction "return"
    // There can be at most one result Parameter per Operation.
    // The parameters of an Exception must all have the direction "out".
    attribute DirectionKind direction;

    // A parameter may have multiple values ...
    attribute MultiplicityType multiplicity;
};

// A Constraint model element is used to express semantic
```

```
// constraints on constrained elements
class QUOTE "Constraint" : ModelElement {
    constraint CannotConstrainThisElement : "[C-50]";
    constraint ConstraintsLimitedToContainer : "[C-51]";

    attribute any expression;

    attribute string language;

    enum EvaluationKind {immediate, deferred};

    attribute EvaluationKind evaluation_policy;

    known constrained_elements projects QUOTE "constraint" of Constrains;
};

// A LiteralType is expressed using the same conventions
// as the CORBA IR uses; e.g. strings are used for enumeration
// values.
typedef any LiteralType;

// A Constant model element describes a binding between a
// a name and a typed value.
class Constant : TypedElement {
    constraint ConstantsValueMustMatchType : "[C-52]";
    constraint ConstantsTypeMustBeSimpleDataType : "[C-53]";

    // The value of a constant value
    attribute LiteralType value;
};

// A Tag is the basis a general mechanism for attaching name/value pairs
// to a model element. Within a model it can be used to attach "pragmas"
// etc that modify the meaning of the model.
class Tag : ModelElement {
    attribute string tag_id;
    attribute bag [0..*] of any values;
    known elements projects tag of AttachesTo;
};
```



```
association AttachesTo {
    role set [1..*] of ModelElement model_element;
    role ordered set [0..*] of Tag tag;
};

// DependsOn is derived from other associations in the model.
// The intended semantic is that ME-1 depends on ME-2 if ME-2 is
// a part of the definition of ME-1. For example, an ME depends
// on its Constraints, a Namespace depends on its contents, and
// Attribute depends on the 'attr_type' Classifier that gives
// its type.
derived association DependsOn {
    readonly role set [0..*] of ModelElement dependent;
    readonly role set [0..*] of ModelElement provider;
};

// Containment is constrained by the "feature matrix". In the case
// of Association <-- Contains --> AssociationEnd, the cardinality is
// constrained as well.
association Contains {
    composite role set [0..1] of Namespace container;
    role ordered set [0..*] of ModelElement contained_element;
};

// Generalizes expressed supertype / subtype relationships between
// Classes and Packages.
//
// We say that the supertype generalizes the subtype.
association Generalizes {
    role ordered set [0..*] of GeneralizableElement supertype;
    role set [0..*] of GeneralizableElement subtype;
};

// This association represents importing of external objects
// into a namespace.
association Aliases {
    role set [0..*] of QUOTE "Import" importer;
    role single Namespace imported;
};
```

```
// This association attaches a constraint to a model element.
// The constraint and model element should belong to the same package.
association Constrains {
    role set [0..*] of QUOTE "Constraint" QUOTE "constraint";
    role set [1..*] of ModelElement constrained_element;
};

// This association lists the exceptions that an operation may raise
association CanRaise {
    role set [0..*] of Operation operation;
    role ordered set [0..*] of QUOTE "Exception" except;
};

// The following two associations relate References to the Ends of
// an Association. The End exposed by a Reference is the one that
// the Reference's container Class fills.
derived association Exposes {
    role set [0..*] of Reference referrer;
    role single AssociationEnd exposed_end;
};

// The role ref'd by an Reference is the one that the Reference
// allows a client to navigate to.
association RefersTo {
    role set [0..*] of Reference referent;
    role single AssociationEnd referenced_end;
};

association IsOfType {
    role single Classifier type;
    role set [0..*] of TypedElement typed_elements;
};

};
```

MOF Implementation Requirements *D*

D.1 Introduction

This specification seeks to avoid any undue implementation requirements, relying on the experience and ingenuity of vendors to exceed any proscribed design. However, to support interoperability of implementations, there are a few places where specific approaches are required. The MOF specification expects interoperability among facilities developed and provided from different vendors. This interoperability includes:

- Model interoperability. The ability to transfer a model developed in one facility to another facility, with no loss or corruption of information.
- Repository interoperability. The ability of a model under development in one facility to import and use Packages and elements of Packages owned by another facility.
- Client interoperability. The ability of tools or other software developed to use one vendor's MOF to make use of another vendor's MOF without change.

D.2 Vendor Boundaries

The implementation requirements are needed to ensure that different vendors take a compatible approach to implementing certain features, when incompatible approaches risk the loss of interoperability. This required compatibility almost always involves object interactions. Yet, it is recognized that the great majority of these object interactions will remain within one vendor's boundary - the interacting objects will all have implementations from the same vendor. Rather than saddle vendors with these implementation requirements unilaterally, they are only required across vendor boundaries. This relaxation of requirements is born from the recognition that these implementation requirements will likely not end up being the optimal implementations.

Determining vendor boundaries will be left to implementations. However, at a minimum, a MofRepository defines the vendor boundary. If an implementation can safely identify a more inclusive boundary, they are welcome to do so. Within an individual MofRepository; however, the implementation is insured to be provided from a single vendor.

D.3 Requirements to Support Associations Across Vendor Boundaries

Due to the nature of composition, it is not possible for a model to contain elements from other Repositories. A Package cannot contain a Package or a Type found in some Repository outside its own. However, through the Import mechanism, a model can make use of most model elements in other Repositories.

Although there are multiple ways in which Associations, References, and Links can be implemented, it is necessary to define a consistent implementation to the object level to ensure interoperability. Because at least one participant in an Association crossing repository boundaries will not have a Reference defined for the Association, the responsibility for maintaining Links falls to the Association.

For any Association which crosses Repository boundaries to another Repository implementation, any invocation of a Reference of that Association will result in a corresponding invocation to the Association itself. For Association X of Type A and Type B, with corresponding AssociationEnds `ae` and `be`, and a reference in Type A of `bref`, use of:

```
a.addBref(b)
```

results in:

```
X.add(a, b).
```

A call of:

```
a.removeBref(b)
```

results in:

```
X.remove(a, b).
```

Calling:

```
a.addBrefBefore(b, b1)
```

results in:

```
X.add_before_be(a, b, b1)
```

The operation:

```
a.bref()
```

must provide the same results as:

```
X.with_ae(a)
```

E.1 Introduction

This appendix summarizes potential areas of future work related to the MOF based on feedback of MOF submitters and reviewers. Note that as with most OMG technologies that are being standardized an abundance of ideas are proposed. Some simple extensions such as the support of higher order associations to more complex ones such as MOF model versioning have been proposed.

Additional work is anticipated in extending the proposed MOF standard to address the related standards such as EIA CDIF and RM-ODP.

E.2 Extending the MOF to Support Ternary and Higher Order Associations

The decision to support only binary associations was based on patterns of association use in industry modeling, the additional encumbrances placed on interfaces when ternary and higher-order associations are introduced, and the additional requirements for completely specifying cardinality constraints (multiplicity). In the tradeoff between simplicity and expressive power, the submitters chose simplicity. The submitters believe that the MOF can be extended in the future to support N-ary associations with minimal impact to current MOF applications.

E.3 Support of Stream based Interchange Format

A stream based interchange mechanism as has been defined by CDIF is a useful mechanism to exchange MOF and UML compliant models - especially for legacy tools which have traditionally not supported programmatic interfaces for exchanging models and model fragments. The MOF designers have anticipated the need for such a mechanism and expect to accommodate this requirement in an upwardly compatible manner. The Package class in the MOF includes internalize and externalize operations to support this capability.

The MOF and OA&DF submitters recommend that a stream based interchange mechanism be the subject of a future OMG RFP.

E.4 Support for MOF Evolution and Versioning

The MOF and MOF Compliant metamodels will evolve over time. The issue of meta model and instance evolution can be solved using a variety of techniques including the use of versioning. This (and related issues like Interface Versioning) are critical issues that needs to be addressed in enterprise development and runtime environments. The MOF submitters recommend that this topic be a subject of a future RFP.

E.5 Support for Mapping between Models

Transformation between metamodels and models is of interest to the tool vendor and end user community to provide interoperability between multiple type systems. The MOF specification defines mappings between MOF and CORBA IDL. A desirable direction for the MOF is to provide a more general purpose framework and a set of interfaces for transformation between meta models. A related topic suggested during the MOF evaluation period is that of generating IDL for "extensions" to the MOF model.

E.6 Interoperability with Microsoft Repository

Microsoft has efforts underway to create a series of COM based repository and information model (metamodel) interfaces in a number of application development technology domains such as object modeling(UML), database management, component management etc. Similar efforts are already underway at the OMG based on the OA&DF, MOF, BOF and the CORBA Component model efforts creating a critical mass of component software enabling standards.

While the OA&DF (UML meta model) specification has 'universal' support (in terms of endorsements from the OMG community and Microsoft), such support does not extend to the corresponding CORBA interfaces. It is also possible that these models may diverge in the future leading to interoperability problems for users of UML and the MOF. A similar problem with COM/CORBA and more recently DCOM/CORBA interoperability has resulted in related OMG RFPs to address the problem.

Future OMG RFPs to address repository and meta model interoperability between CORBA and DCOM environments is crucial for customers and vendors who have invested in both the technologies. Of course if Microsoft technologies supported CORBA based information models as well (as is partially the case with UML), such an RFP would be unnecessary and the user community would have consistent information models, components, and compatible tools.

-
- A**
- addLink 6-25
 - addLinkBefore 6-25
 - addValue 6-13
 - addValueAt 6-15
 - addValueBefore 6-14
 - aggregation 3-50
 - AggregationType 3-80
 - all_links 6-23
 - allObjects 5-58, 5-62, 5-70, 5-72, 6-11
 - allSupertypes 3-26
 - annotation 3-15
 - Annotation Template 5-92
 - AnnotationType 3-77
 - Architecture, four layer metamodel 2-2
 - Association 3-47
 - Association Template 5-59
 - Associations 3-8
 - Attribute Template 5-68
 - Attributes 3-4
- B**
- BehavioralFeature 3-43
- C**
- Class 3-31
 - Collection Kinds 5-42
 - Common Exceptions 5-24
 - Complex bindings 1-5
 - Constant Template 5-89
 - constrainedElement 3-73
 - constrainedElements 3-60
 - Constraint 3-59
 - constraint 3-73
 - Constraint Template 5-91
 - Constraint-Constrains-ModelElement 3-73
 - Constraints 3-16
 - Contained Elements 3-4
 - containedElement 3-66
 - container 3-16, 3-65
 - contents 3-22
 - CORBA
 - documentation set i
 - CORBA Basic Types 3-77
 - CORBA IDL for the Meta Object Facility B-1
 - CORBA IDL for the Reflective Interfaces 6-30
 - createInstance 6-10
- D**
- Data Type Template 5-90
 - Data Types 6-30
 - Data viewpoint 1-1
 - Data warehouse management scenarios 1-6
 - delete 6-8
 - DependencyKind 3-81
 - dependent 3-74
 - DepthType 3-79
 - direction 3-58
 - DirectionType 3-80
- E**
- elements 3-64
 - Ends 3-8, 3-65
 - evaluationPolicy 3-60
 - EvaluationType 3-80
 - exception 3-71
 - Exception Template 5-89
 - exceptions 3-45
 - exposedEnd 3-42, 3-69
 - expression 3-59
 - Extending the MOF to Support Ternary and Higher Order Associations E-1
 - externalize 3-54
- F**
- Feature 3-36
 - findElementsByType 3-23
 - findElementsByTypeExtended 3-27
 - findRequiredElement 3-18
 - Format 1 5-41
 - Format 2 5-41
 - FormatNotSupported 3-84
 - Future Directions for the MOF E-1
- G**
- GeneralizableElement type 3-24
 - GeneralizableElement-Generalizes-GeneralizableElement 3-66
 - Generation Rules for Collection Kinds 5-42
 - getAssociation 6-29
 - getClassRef 6-28
 - getNestedPackage 6-29
- I**
- Identifier Format 1 5-41
 - Identifier Format 2 5-41
 - Identifier Name Scoping 5-44
 - Identifier Naming 5-39
 - IDL for the Reflective Interfaces 6-30
 - IDL mapping 5-1
 - immediate_containing_package 6-7
 - Import 3-55
 - Import-Aliases-Namespace 3-72
 - imported 3-72
 - importedNamespace 3-56
 - importer 3-72
 - Information management scenarios 1-6
 - Interface Repository (IR) 1-4
 - internalize 3-54
 - Interoperability with Microsoft Repository E-2
 - interpreting IDL templates 5-46
 - invokeOperation 6-20
 - isAbstract 3-25
 - isChangeable 3-39, 3-51
 - isDerived 3-40, 3-48
 - isFrozen 3-17
 - isInstanceOf 6-10
 - isLeaf 3-26, 3-78, 3-82, 3-83
 - isNavigable 3-50
 - isQuery 3-44
 - isRequiredBecause 3-19
 - isRoot 3-25
 - isSingleton 3-32
 - isVisible 3-18

Index

L

language 3-60
link_exists 6-24
Literal String Values 5-42
lookupElement 3-22
lookupElementExtended 3-27

M

mapping 5-1
Mapping Rules 5-45
Metamodel architecture 2-2
metaObject 6-7
ModelElement 3-15
modelElement 3-76
ModelElement Containment Rules 3-12
ModelElement-DependsOn-ModelElement 3-74
Modeling viewpoint 1-1
modifyLink 6-26
modifyValue 6-15, 6-16
modifyValueAt 6-16, 6-17
MODL Description of the MOF C-1
MOF Implementation Requirements D-1
MOF Model C-1
MOF Model Associations 3-65
MOF Model Data Types 3-77
MOF Model Exceptions 3-83
MOF Model IDL A-1, B-1
MOF model types 3-3, 3-15
MOF to IDL Mapping 4-1, 5-1
MofAttribute 3-40
MofException 3-46
multiplicity 3-39, 3-50, 3-58, 5-52, 5-53, 5-56
MultiplicityType 3-78, 3-82

N

name 3-15
nameIsValid 3-23
NameNotResolved 3-83
Namespace type 3-21
Namespace-Contains-ModelElement 3-65
NameType 3-77
Notation 5-46

O

Object Management Group i
 address of ii
ObjectNotExternalizable 3-84
OCL Representation of the MOF Model Constraints 3-86
OMG Architecture and the Repository Common Facility ix
operation 3-71
Operation Template 5-86
Operation-CanRaise-MofException 3-71
Operations 3-7
otherEnd 3-51
outermost_container 6-20
outermost_containing_package 6-8

P

Package 3-53
Package Create Template 5-49
Package Template mapping rules 5-46
Preconditions for Successful IDL mapping 5-33

provider 3-75

Q

qualified Name 3-15
query 6-24

R

Reference 3-41
referencedEnd 3-41
referenceEnd 3-68
Reference-Exposes-AssociationEnd derived 3-69
Reference-RefersTo-AssociationEnd 3-68
References 3-5
referent 3-68
referrer 3-69
reflItself 6-7
Reflection IDL B-22
Reflective
 RefAssociation 6-22
 RefBaseObject 6-5
 RefPackage 6-28
Reflective Exceptions 5-31
Reflective Module 6-3
removeLink 6-27
removeValue 6-17, 6-18
removeValueAt 6-18
Repository Common Facility ix
repository service 1-3
requiredElements 3-16
Requirements to Support Associations Across Vendor
 Boundaries D-2
resolveQualifiedName 3-22
Rules 5-45
Rules for Splitting MOF Model Names into "Words" 5-40
Rules of ModelElement Containment 3-12

S

scope 3-38
ScopeType 3-80
Service interface bridges 1-5
setValue 6-12, 6-13
Software development scenarios 1-2
Stereotypes iii, C-1, D-1, E-1
StructuralFeature 3-38
subtype 3-67
Successful IDL mapping 5-33
supertype 3-67
Supertypes 3-4
supertypes 3-26
Support for Mapping between Models E-2
Support for MOF Evolution and Versioning E-2
Support of Stream based Interchange Format E-1

T

Tag 3-63
tag 3-76
Tag-AttachesTo-ModelElement 3-76
tagId 3-64
type 3-29, 3-70
Type Create Template 5-57
Type Forward Declaration Template 5-51

Type management scenarios 1-4
Type Template 5-54
TypeAlias 3-35
typeCode 3-34
TypedElement type 3-29
TypedElement-IsOfType-Classifier 3-70
typedElements 3-70
TypeDescriptor 3-77
Types 3-3

U

UDL development system 1-4

usage scenario for repository service 1-3

V

value 3-62, 6-11
values 3-64
Vendor Boundaries D-1
VerificationResultKind 3-81, 3-82
verify 3-17
visibility 3-25, 3-37, 3-55
VisibilityType 3-79

