Lecture Notes in Computer Science2900Edited by G. Goos, J. Hartmanis, and J. van Leeuwen

Springer Berlin

Berlin Heidelberg New York Hong Kong London Milan Paris Tokyo Michel Bidoit Peter D. Mosses

CASL User Manual

Introduction to Using the Common Algebraic Specification Language

With chapters by

Till Mossakowski, Donald Sannella, and Andrzej Tarlecki



Series Editors

Gerhard Goos, Karlsruhe University, Germany Juris Hartmanis, Cornell University, NY, USA Jan van Leeuwen, Utrecht University, The Netherlands

Authors

Michel Bidoit Laboratoire Spécification et Vérification, CNRS UMR 8643 École Normale Supérieure de Cachan 61, Avenue du Président Wilson, 94235 Cachan Cedex, France E-mail: bidoit@lsv.ens-cachan.fr

Peter D. Mosses University of Aarhus, BRICS and Department of Computer Science Aabogade 34, 8200 Aarhus N, Denmark E-mail: pdmosses@brics.dk

Cataloging-in-Publication Data applied for

A catalog record for this book is available from the Library of Congress.

Bibliographic information published by Die Deutsche Bibliothek Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data is available in the Internet at <http://dnb.ddb.de>.

CR Subject Classification (1998): D.2.1, D.3.1, D.2, D.3, F.3

ISSN 0302-9743 ISBN 3-540-20766-X Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag is a part of Springer Science+Business Media

springeronline.com

© 2004 IFIP International Federation for Information Processing, Hofstrasse 3, A-2361 Laxenburg, Austria Printed in Germany

Typesetting: Camera-ready by author, data conversion by Olgun Computergrafik Printed on acid-free paper SPIN: 10931011 06/3142 5 4 3 2 1 0

Preface

CASL, the *Common Algebraic Specification Language*, has been designed by CoFI, the *Common Framework Initiative* for algebraic specification and development. CASL is an expressive language for specifying requirements and design for conventional software. It is algebraic in the sense that models of CASL specifications are algebras; the axioms can be arbitrary first-order formulas.

This User Manual illustrates and discusses how to write CASL specifications.

CASL is a major new algebraic specification language. It has been carefully designed by a large group of experts as a general-purpose language for practical use in software development – in particular, for specifying both requirements and design. CASL includes carefully-selected features from many previous specification languages, as well as some novel features that allow algebraic specifications to be written much more concisely and perspicuously than hitherto. It may ultimately replace most of the previous languages, and provide a common basis for future research and development.

CASL has already attracted widespread interest within the algebraic specification community, and is generally regarded as a de facto standard. Various sublanguages of CASL are available – primarily for use in connection with existing tools that were developed in connection with previous languages. Extensions of CASL provide languages oriented toward development of particular kinds of software (reactive, concurrent, etc.).

Major libraries of validated CASL specifications are freely available on the Internet, and the specifications can be reused simply by referring to their names. Tools are provided to support practical use of CASL: checking the correctness of specifications, proving facts about them, and managing the formal software development process.

The companion CASL *Reference Manual* [20] provides full details of the CASL design, including its formal semantics.

After briefly reviewing the background of CoFI and CASL, and the underlying concepts of algebraic specification languages, this book introduces the potential user to the features of CASL mainly by means of illustrative examples. It presents and discusses the typical ways in which the language concepts and constructs are expected to be used in the course of building system specifications. Thus, the presentation focuses on *what* the constructs and concepts of CASL are *for*, and *how* they should (and should not) be used. These points are made as clear as possible by referring to simple examples, and by discussing both the general ideas and some details of CASL specifications.

Further chapters introduce the reader to the CASL Reference Manual, to some of the currently available CASL support tools, and to a couple of the CASL libraries of basic datatypes. A substantial case study of the practical use of CASL in an industrially-relevant context completes the material. Appendices provide a quick reference of CASL constructs, a list of the main points to bear in mind when using CASL, and the original informal requirements for the case study.

Structure

Part I: Background

Chapter 1 describes the origins of CASL: how CoFI was formed in response to the proliferation of algebraic specification languages in the preceding two decades, and the aims and scope that were formulated for this international initiative.

For the benefit of readers not already familiar with other algebraic specification languages, Chap. 2 reviews the main concepts of algebraic specification, explaining standard terminology regarding specification language constructs and models (i.e., algebras).

Part II: Writing CASL Specifications

Chapter 3 shows how some familiar datatypes involving total functions are specified in CASL, essentially as in many other algebraic specification languages. Loose, generated, and free specifications are discussed in turn, with illustrative examples and advice on the use of different specification styles.

Partial functions arise naturally. Chapter 4 explains how CASL supports specification of partial functions, drawing attention to where special care is needed compared to specifications involving only total functions.

Subsorts and supersorts are often useful in CASL specifications. Chapter 5 illustrates how they can be declared and defined, and that they can sometimes be used to avoid the need for partial functions.

The examples given so far make use of named and structured specifications in a simple and natural way. Chapter 6 takes a much closer look at the constructs CASL provides for structuring specifications, explaining how large and complex specifications are easily built out of simpler ones by means of a small number of specification-building operations.

Chapter 7 shows how making a specification generic (when appropriate) improves its reusability, allowing it to be instantiated with different arguments; compound identifiers avoid the need for explicit renaming when combining the results of different instantiations. It also introduces the constructs for expressing so-called views between specifications.

While specification-building operations are useful to structure the text of large specifications, architectural specifications are meant for imposing structure on implementations. Chapter 8 discusses and illustrates the role of architectural specifications, and shows how to express them in CASL.

Chapter 9 explains and illustrates how libraries of named specifications can be formed, and made available over the Internet, to encourage widespread reuse and evolution of specifications. Version control is of crucial importance here.

Part III: Carrying On

Chapter 10 gives a detailed overview of the foundations of CASL, which are established in the accompanying CASL Reference Manual.

Tool support is vital for efficient use of formal specifications in connection with practical software design and development. Chapter 11 presents the main tools that have been implemented so far; several of them allow use of CASL specifications in connection with tools that were originally developed for other specification languages, showing how CASL provides tool interoperability.

Chapter 12 introduces a few of the many specifications that are available in the CASL libraries of basic datatypes.

Finally, Chap. 13 gives a realistic case-study of the use of CASL in practice, in connection with the design of software for a Steam-Boiler Control System. This particular example is one of the standard bench-marks for comparing specification frameworks [1].

Appendices and Indexes

This volume is completed by three appendices: App. A provides a compact overview of all CASL constructs, for quick reference; App. B lists all the main points to bear in mind when using CASL; and App. C reproduces the informal requirements specification for the case study.

The names of all the specifications given in this book are listed at the back, together with an index of concepts and a list of references to the literature. (A comprehensive annotated bibliography of publications involving CASL is provided in the Reference Manual.)

An accompanying CD-ROM contains the source files for all the illustrative specifications, and a copy of the libraries of specifications of basic datatypes.

Organization

All the main points are highlighted like this.

The material in this book is organized in a tutorial fashion. Each main point is usually accompanied by an illustrative example of a complete CASL specification; the names of these specifications are listed (both in order of presentation and alphabetically) at the end of the book. Moreover, the points themselves are repeated (in order of presentation) in App. B.

Readers who are familiar with previous algebraic specification languages, and especially those who have been following or participating in the design and development of CASL, may prefer to skip lightly through Chaps. 1 and 2. Chapter 3, however, is mandatory, since it is there that many CASL features needed to understand the subsequent chapters are introduced.

In contrast, Chaps. 4 and 5 can be skipped at first reading if the reader is not so much interested in partial functions, resp. subsorting (with the proviso though that there are some references to the examples given in these chapters from later chapters).

Chapters 6 and 7 present mainstream material, and until one feels comfortable with all the main points and examples, it is advisable to wait with proceeding to Chaps. 8 and 9.

Chapter 10 is primarily for those who will want to follow up on this book with a more detailed study of CASL, based on the Reference Manual. Part of Chap. 11 assumes familiarity with concepts introduced in Chap. 7. In Chap. 12, Sect. 12.1 assumes Chaps. 4 and 5, whereas Sect. 12.2 assumes also Chaps. 6 and 7. Finally, most of Chap. 13 can be studied after Chap. 7, but Sect. 13.10 requires Chap. 8.

Acknowledgement. Chapter 10 was written by Donald Sannella and Andrzej Tarlecki. Chapter 11 was written by Till Mossakowski, with contributions from Mark van den Brand and Markus Roggenbach. Till Mossakowski also provided Chap. 12, based on libraries of CASL specifications developed at the Bremen Institute for Secure Systems in joint work with Markus Roggenbach and Lutz Schröder; and he checked the well-formedness of all the specifications in this book, using CATS, the CASL Tool Set.

Public drafts of this book were released in July and October 2003. The many insightful comments from CoFI participants and other readers were very helpful to the authors during the preparation of the final version. Detailed comments on all or part of the public drafts were received from Jørgen Iversen, Christian Maeder, Guillem Marpons, Narciso Martí-Oliet, Till Mossakowski, Don Sannella, Giuseppe Scollo, Andrzej Tarlecki, Frédéric Voisin, and Alexandre Zamulin. Responsibility for any mistakes in the final version belongs, of course, to the authors.

Michel Bidoit gratefully acknowledges support from LSV,¹ CNRS, and École normale supérieure de Cachan. Peter Mosses gratefully acknowledges support from BRICS² and the Department of Computer Science, University of Aarhus. Much of the material on which this book is based was developed in connection with activities of CoFI-WG (ESPRIT Working Group 29432) and IFIP WG 1.3 (Foundations of System Specification).

Finally, special thanks to Springer, and in particular to Alfred Hofmann as Executive Editor, for their willingness to publish this book, and for helpful advice concerning its preparation.

Michel Bidoit and Peter D. Mosses October, 2003

¹ Laboratoire Spécification et Vérification (www.lsv.ens-cachan.fr).

² Basic Research in Computer Science (www.brics.dk), funded by the Danish National Research Foundation.

Contents

Part I Background

1	Intro	oduction	3
	1.1	CoFI	3
	1.2	Casl	7
2	Und	erlying Concepts	11
	2.1	Basic Specifications	11
	2.2	Structured Specifications	15
	2.3	Architectural Specifications	19
	2.4	Libraries of Specifications	20

Part II CASL Specifications

3	Gett 3.1 3.2 3.3	ing Started Loose Specifications Generated Specifications Free Specifications	23 24 33 36
4	Parti	ial Functions	47
	4.1	Declaring Partial Functions	47
	4.2	Specifying Domains of Definition	50
	4.3	Partial Selectors and Constructors	54
	4.4	Existential Equality	55
5	Subs	orting	57
	5.1	Subsort Declarations and Definitions	57
	5.2	Subsorts and Overloading	61
	5.3	Subsorts and Partiality	62

6	Structuring Specifications		
	6.1	Union and Extension	i7
	6.2	Renaming	;9
	6.3	Hiding	'1
	6.4	Local Specifications	73
	6.5	Named Specifications	'5
7	Gene	eric Specifications	7
	7.1	Parameters and Instantiation	78
	7.2	Compound Symbols 8	35
	7.3	Generic Specifications with Imports	38
	7.4	Views)0
8	Spec	ifying the Architecture of Implementations)3
	8.1	Architectural Specifications)5
	8.2	Generic Components)()
	8.3	Writing Meaningful Architectural Specifications10)6
9	Libra	aries	1
	9.1	Local Libraries	2
	9.2	Distributed Libraries	6
	9.3	Version Control	20

Part III Carrying On

10	Foundations	
11	Tools	s
	11.1	The Heterogeneous Tool Set (HETS)
	11.2	Hol-Casl
	11.3	AsF+SDF Parser and Syntax-Directed Editor
	11.4	Other Tools
12	Basic	c Libraries
	12.1	Library Basic/Numbers
	12.2	Library BASIC/STRUCTUREDDATATYPES
13	Case	Study: The Steam-Boiler Control System155
	13.1	Introduction
	13.2	Getting Started
	13.3	Carrying On161
	13.4	Specifying the Mode of Operation163
	13.5	Specifying the Detection of Equipment Failures
	13.6	Predicting the Behavior of the Steam-Boiler
	13.7	Specifying the Messages to Send

13.8	The Steam-Boiler Control System Specification	183
13.9	Validation of the CASL Requirements Specification	184
13.10	Designing the Architecture	186

Appendices

Α	Casl A.1 A.2 A.3 A.4	Quick Reference193Basic Specifications194Structured Specifications199Architectural Specifications200Libraries201
В	Point B.1 B.2 B.3 B.4 B.5 B.6 B.7 B.8 B.9 B.10 B.11 B.12	ts to Bear in Mind203Introduction203Underlying Concepts203Getting Started204Partial Functions205Subsorting206Structuring Specifications207Generic Specifications207Specifying the Architecture of Implementations207Libraries208Foundations209Tools209Basic Libraries210
С	The C.1 C.2 C.3 C.4 C.5 C.6 C.7	Steam-Boiler Control Specification Problem.211Introduction.211Physical Environment.212The Overall Operation of the Program.214Operation Modes of the Program.214Messages Sent by the Program.216Messages Received by the Program.217Detection of Equipment Failures.219
Ref	ferenc	es
List	t of N	amed Specifications
Ind	ex of	Library and Specification Names
Сол	ncept	Index

Introduction

This chapter first explains the background and aims of CoFI, the Common Framework Initiative for algebraic specification and development of software. It then gives an overview of the main features of CASL, the Common Algebraic Specification Language.

1.1 CoFI

In 1995, an open collaborative effort was initiated: to design a common framework for algebraic specification and development of software. It is referred to as *The Common Framework Initiative*, $CoFI.^1$

There was an urgent need for a common framework.

The rationale behind this initiative was that the lack of such a common framework was a major hindrance for the dissemination and application of algebraic specification techniques. In particular, there was a proliferation of languages – some differing in only quite minor ways from each other. The major languages included ACT ONE/ACT TWO [19], ASF [6], ASL [36], CLEAR [15], EXTENDED ML [26], LARCH [25], OBJ3 [24], PLUSS [9], and SPECTRUM [14]. This abundance of languages was an obstacle for the adoption of algebraic methods for use in industrial contexts, making it difficult to exploit standard examples, case studies and training material. A common framework, with widespread support at least throughout the research community, was urgently needed.

¹ CoFI is pronounced like 'coffee'.

[©] IFIP International Federation for Information Processing 2004

The aim of CoFI was to base the common framework as much as possible on a critical selection of features that had already been explored in previous research and applications (see the IFIP State-of-the-Art Report on *Algebraic Foundations of Systems Specification* [3] for the background). The collective experience and expertise of the CoFI participants provided a unique opportunity to achieve this aim within a reasonably short time-span.

The various groups working on algebraic specification frameworks had already had ample opportunity to develop their own particular variations on the theme of algebraic specification [16], yet no clear 'winner' had emerged (although there were several strong contenders).

CoFI aims at establishing a wide consensus.

The aim of CoFI was to design a framework incorporating just those features for which there would be a wide consensus regarding their appropriateness. This framework should be able to subsume many of the existing frameworks, and be seen as an attractive common basis for future research and development – with high potential for strong collaboration between the various groups.

The initial overall aims of CoFI were formulated as follows:

- A common framework for algebraic specification and software development is to be designed, developed, and disseminated.
- The production of the common framework is to be a collaborative effort, involving a large number of experts (30–50) from many different groups (20–30) working on algebraic specification.
- In the short term, the common framework is to become accepted as an appropriate basis for a significant proportion of the research and development in algebraic specification.
- Specifications in the common framework are to have a uniform, userfriendly syntax and straightforward semantics.
- The common framework is to be able to replace many existing algebraic specification frameworks.
- The common framework is to be supported by a concise reference manual, user's guide, libraries of specifications, tools, and educational materials.
- In the longer term, the common framework is to be made attractive for use in industrial contexts.
- The common framework is to be available free of charge, both to academic institutions and to industrial companies. It is to be protected against appropriation.

The focus of CoFI is on algebraic techniques.

The functionality of the common framework is to allow and be useful for:

- algebraic specification of the functional requirements of software systems, for some significant class of software systems;
- formal development of design specifications from requirements specifications, using some particular methods;
- documenting the relation between informal statements of requirements and formal specifications;
- verification of correctness of development steps from (formal) requirements to design specifications;
- documenting the relation between design specifications and implementations in software;
- exploration of the (logical) consequences of specifications: e.g., rewriting, theorem-proving, prototyping;
- reuse of parts of specifications;
- adjustment of specifications and developments to changes in requirements;
- providing a library of useful specification modules; and
- providing a workbench of tools supporting the above.

CoFI has already achieved its main aims.

The first major achievement of CoFI was the completion of the design of CASL, the Common Algebraic Specification Language. The CASL design effort started in September 1995. An initial design was proposed in May 1997 (with a language summary, abstract syntax, and formal semantics, but no concrete syntax) and tentatively approved by IFIP WG1.3. The report of the IFIP referees on the initial CASL design proposal suggested reconsideration of several points in the language design, and requested some improvements to the documents describing the design. Apart from a few details, the design was finalized in April 1998, and CASL version 1.0 was released in October 1998. IFIP WG 1.3 was asked to review the final design of CASL version 1.0.1 in May 2000, and subsequently approved that design in April 2001. The current version (1.0.2) was adopted in October 2003; it incorporates adjustments to some minor details of the concrete syntax and the semantics. No further revisions of the CASL design are anticipated.

The CASL *Reference Manual* [20], published as a companion volume to the present book, includes a detailed yet concise (60 pages) summary of the CASL design; the rest of it is concerned mainly with the formal syntax and semantics of CASL, and with the libraries of basic datatype specifications. An introduction to the Reference Manual is given in Chap. 10 of this book. In parallel with the design of CASL, CoFI has developed tool support for the use of CASL (see Chap. 11), and substantial libraries of CASL specifications (see Chap. 12).

Despite the previous lack of a CASL User Manual, there is already much evidence that CASL is now accepted as an appropriate basis for research and development in algebraic specification. Reference [33] gives an overview of what was achieved in the period 1998–2001, and the annotated bibliography in the CASL Reference Manual lists a significant number of further publications that involve CASL. At the time of writing, it remains to be seen whether significant industrial take-up will follow.

CoFI is an open, voluntary initiative.

CoFI was started by COMPASS (ESPRIT Basic Research WG 3264/6112, 1989-96), in cooperation with IFIP WG 1.3 (Foundations of System Specification, founded 1992), on the basis of proposals made during their meetings in 1994 (Santa Margherita Ligure, Italy) and 1995 (Oslo, Norway); participation in CoFI was, however, never confined to members of those working groups. The active participants have included some 30 leading researchers in algebraic specification, with representatives from almost all the European groups working in this area. (Ideally, representatives from non-European groups would have been involved too, but logistic difficulties prevented this.)

Originally, CoFI had separate task groups concerning language design, semantics, tools, methodology, and reactive systems. There was a substantial amount of interaction between the task groups, which was facilitated by many of the CoFI participants being involved in more than one task group. The overall coordination of these task groups was managed by Peter Mosses from the start of CoFI in September 1995 until August 1998, and subsequently by Don Sannella. In 2003, the CoFI task groups were replaced by a looser coordination mechanism, with a steering committee chaired by Don Sannella.

CoFI has received funding as an ESPRIT Working Group, and is sponsored by IFIP WG 1.3.

The European Commission provided funding for the European component of CoFI as ESPRIT Working Group 29432 from 1998 to 2001 [33]. The partners were the coordinating sites of the various CoFI task groups (University of Aarhus, University of Bremen, École normale supérieure de Cachan, University of Genova, INRIA Lorraine, Warsaw University) with the University of Edinburgh as overall coordinator. The goals of the working group were to coordinate the completion of and disseminate the Common Framework, to demonstrate its practical applicability in industrial contexts, and to establish the infrastructure needed for future European collaborative research in algebraic techniques. Apart from this period of funding, and support for meetings from the COMPASS Working Group until its termination in 1996, CoFI has relied entirely on unfunded efforts by its participants. Participation in the frequent working meetings was often supported by generous subsidies from the local organizers.

IFIP WG 1.3 sponsors CoFI by reviewing proposals for changes to the design of CASL, and proposals for extensions of CASL. Moreover, a considerable number of members of IFIP WG 1.3 have been (and, at the time of writing, still are) active participants of CoFI.

New participants are welcome!

CoFI is an open collaboration, and new participants are always welcome. Current information about CoFI activities is available at the main CoFI web site: http://www.cofi.info. The low-volume mailing list cofi@cofi.info is reserved for CoFI announcements, and discussions about CoFI activities generally take place on the mailing list cofi-discuss@cofi.info; see the CoFI web site for how to subscribe, and for access to the archives.

$1.2 \,\, \mathrm{Casl}$

CASL has been designed as a general-purpose algebraic specification language, subsuming many existing languages.

The primary specification language developed by CoFI is called CASL: the Common Algebraic Specification Language. Its main features are:

- The design of CASL is based on a critical selection of the concepts and constructs found in existing algebraic specification frameworks.
- CASL is an expressive specification language with simple semantics and good pragmatics.
- CASL is appropriate for specifying requirements and design of conventional software packages.
- CASL is at the heart of a coherent family of languages that are obtained as sublanguages or extensions of CASL.

CASL subsumes many previous languages for the formal specification of functional requirements and modular software design. Tools for CASL are interoperable, i.e., capable of being used in combination rather than in isolation. CASL interfaces to existing tools extend this inter-operability (see Chap. 11). 8 1 Introduction

The intention was to base the design of CASL on a critical selection of concepts and constructs from existing specification languages. However, it was not easy to reach a consensus on a coherent language design. A great deal of careful consideration was given to the effect that the constructs available in the language would have on such aspects as the methodology and tools. A complete formal semantics for CASL was produced in parallel with the later stages of the language design, and the desire for a relatively straightforward semantics was one factor in the choice between various alternatives in the design.

CASL represents a consolidation of past work on the design of algebraic specification languages. With a few minor exceptions, all its features are present in some form in other languages, but there is no language that comes close to subsuming it. Designing a language with this particular novel collection of features required solutions to a number of subtle problems in the interaction between features. An overview of the CASL design is presented in [2], and full details are provided in the CASL Reference Manual [20].

CASL is at the center of a family of languages.

It was clear from the start that no single language could suit all purposes. On the one hand, sophisticated features are required to deal with specific programming paradigms and special applications. On the other hand, important methods for prototyping and reasoning about specifications only work in the *absence* of certain features: for instance, term rewriting requires specifications with equational or conditional equational axioms.



Fig. 1.1. The CASL Family of Languages

CASL is therefore at the center of a *family* of languages, see Fig. 1.1. Some tools will make use of well-delineated *sublanguages* of CASL, obtained by syntactic or semantic restrictions [29], while *extensions* of CASL are generally designed to support various paradigms and applications. The design of CASL took into account the need to define sublanguages and extensions.

CASL itself has several major parts.

The major parts of CASL are concerned with *basic* specifications, *structured* specifications, *architectural* specifications, and *libraries* of specifications. They have been designed to be used together: basic specifications can be used in structured specifications, which in turn can be used in architectural specifications; structured and/or architectural specifications can be collected into libraries. However, these parts of CASL are quite independent, and may be understood separately, as we shall see in Part II.

Underlying Concepts

CASL is based on standard concepts of algebraic specification.

This chapter reviews the main concepts of algebraic specification. It briefly explains and illustrates standard terminology regarding specification language constructs and models of specifications (i.e., algebras), and indicates the differences between basic, structured, and architectural specifications.

The focus here is on concepts that are relevant to CASL, and which will be needed in later chapters. For comprehensive presentations of concepts and results concerning algebraic specification, see [3, 10, 16, 27, 34, 35, 37]; for an overview of the design of CASL, see [2]; and for full details of CASL, see the CASL Reference Manual [20].

The reader is assumed to be familiar with basic mathematical notions (sets, relations, and total and partial functions) and with the use of logical formulas as axioms.

2.1 Basic Specifications

A basic specification declares symbols, and gives axioms and constraints.

A basic specification in an algebraic specification language generally consists of a set of *declarations* of *symbols*, and a set of *axioms* and *constraints*, which restrict the *interpretations* of the declared symbols. CASL allows basic specifications to include also items which simultaneously declare symbols and restrict their interpretations. The semantics of a basic specification is a signature and a class of models.

The meaning or semantics of a basic specification SP generally has two parts:

- a signature Σ , corresponding to the symbols introduced by the specification, and
- a class of Σ -models,¹ corresponding to those interpretations of the signature Σ that satisfy the axioms and constraints of the specification.

When a model M satisfies a specification SP, we write $M \models SP$ and say that M is a model of SP. (Formalizing this within the theory of so-called *institutions* involves categorical structure on the set of signatures and on the class of models, and a natural condition on the *satisfaction relation*. We need not bother with the details here – but see however the concept of a signature morphism in Sect. 2.2.) A specification is said to be *consistent* when its class of models is non-empty, and otherwise *inconsistent*.

CASL specifications may declare sorts, subsorts, operations, and predicates.

A CASL *signature* represents declarations of sorts, subsorts, operations, and predicates. The signature is called *many-sorted* when there are no subsort declarations, and otherwise *subsorted*; it is called *algebraic* when there are no predicate declarations.

Sorts are interpreted as carrier sets.

A sort is a symbol which is interpreted as a set, called a *carrier set*. The elements of a carrier set are generally abstract representations of the data processed by software: numbers, characters, lists, etc. Thus a sort declared by a specification corresponds to a type in a programming language. Sort symbols are usually chosen to be strongly suggestive of their intended interpretations, e.g., *Int* for a sort to be interpreted as the set of integers, *List* for a set of lists. CASL allows also *compound sort* symbols, such as *List*[*Int*] for lists of integers.

¹ Readers who are not interested in foundational aspects may treat the word 'class' as a synonym for 'set'. In general, the models of an algebraic specification in CASL constitute a *proper class*, because there is no restriction on the elements of the carrier sets.

Subsorts declarations are interpreted as embeddings.

A sort may be declared to be a *subsort* or a *supersort* of other sorts. The subsort relation between two sorts could be interpreted as *set inclusion*. Its interpretation in CASL is however more general: it is interpreted as an *embedding*, i.e., a 1-1 function from the carrier set of the subsort to that of the supersort. For example, if *ASCII* is specified to be a subsort of *ISO_Latin1* in CASL, the carrier set for *ASCII* could be simply a subset of that for *ISO_Latin1*. If *Char* were to be declared as a subsort of *String*, however, the carrier sets for *Char* and *String* could be disjoint, with the embedding mapping each character to the corresponding single-character string. (See also the concept of overloading, below.)

Operations may be declared as total or partial.

An operation symbol consists of the name of the operation together with its *profile*, which indicates the number and sorts of the arguments, and the *result* sort. In CASL, a declared operation symbol is interpreted as either a *total* or a *partial function* from the Cartesian product of the carrier sets of the argument sorts to the carrier set of the result sort; the subset of the argument tuples for which the result of a function is defined is called its *domain of definition*. The declaration indicates whether the function is total or partial.² For example, integer addition would be declared as total, but integer division as partial. The result of applying an operation is undefined whenever any of its arguments is undefined (regardless of whether the operation itself is total or partial).

When there are no arguments, the operation is called a *constant*. A constant is interpreted simply as an element of the result sort.

Predicates are different from boolean-valued operations.

A *predicate* symbol consists of the name of the predicate together with its *profile*, which indicates the number and sorts of the arguments but no result sort: predicates are syntactically different from boolean-valued operations, and are used to form atomic formulas rather than terms. In CASL, a declared predicate symbol is interpreted as a relation on (i.e., a subset of) the Cartesian product of the carrier sets of the argument sorts. An application of a predicate is said to *hold* when the tuple of arguments is in the relation. For example, a

² A partial function might just happen to be everywhere defined, of course.

symbol '<' to be interpreted as the less-than relation could be declared as a binary predicate on integers.

An application of a predicate simply fails to hold when any of its arguments is undefined: there is no undefinedness about holding or not. This allows the logic to remain two-valued, and the logical connectives to have their familiar interpretations.

In contrast, the result of evaluating an application of even a total booleanvalued operation could be true, false, or undefined: the last case arises when any argument of the application is undefined. Thus boolean-valued operations corresponding to logical connectives (conjunction, implication, etc.) have to take account of undefinedness, which leads to a three-valued logic.

A further significant difference between predicates and boolean-valued operations shows up in connection with the concept of initiality, see Sect. 2.2. (Predicates of two-valued logic can be represented accurately by *partial* operations with a *single-valued* result sort, holding being represented by definedness.)

Operation symbols and predicate symbols may be overloaded.

An operation or predicate name can be declared with different profiles in the same specification. This is called *overloading*. For example, the constant '*empty*' could be overloaded, being interpreted as (unrelated) elements of the sorts *List* and *Set*, according to the context of its use. Similarly, a predicate name such as '<' could be overloaded on unrelated sorts such as *Char* and *Int*.

In CASL, overloading is required to be *compatible* with embeddings between subsorts. For example, the sort Nat, interpreted as the set of natural numbers, might be a subsort of Int, interpreted as the set of all integers; then when the operation name '+' and the predicate name '<' are declared both on Nat and on Int, their interpretations are required to be such that it makes no difference whether the embedding from Nat to Int is applied to the arguments or to the result of the operation, and whether it is applied to the arguments of the predicate or not.

Axioms are formulas of first-order logic.

The interpretation of quantification (universal, existential, and uniqueexistential) and of the usual logical connectives (negation, conjunction, disjunction, implication, and equivalence) in CASL axioms is completely standard. Variables in formulas range over the carrier sets of specified sorts. Apart from the usual predicate applications, the atomic formulas in CASL axioms are equations (strong or existential), definedness assertions, and subsort membership assertions. An *existential* equation holds when the values of its terms are defined and equal; a *strong* equation holds moreover when the values of the terms are both undefined.

Regardless of whether the values of the terms occurring in an axiom are defined, the axiom either holds or it does not hold in a particular model: the logic is two-valued, there is no "maybe" or undefinedness about the holding of axioms. Recall that when the value of any argument term is undefined, an application of a predicate never holds; similarly, definedness and subsort membership assertions never hold when their arguments are undefined.

Sort generation constraints eliminate 'junk' from specific carrier sets.

In general, the carrier sets of the models of a specification may contain 'junk' elements, i.e., elements which cannot be obtained by any composition of the operations declared by the signature of the specification.

A sort generation constraint in CASL concerns specific sorts and operations, and is satisfied in a model when no elements of the indicated carrier sets are junk with respect to the indicated operations – i.e., all the elements of those sets can be obtained by consecutively applying those operations to elements of the carrier sets of the remaining sorts. For example, the carrier set for the sort *Container* might be constrained to be generated from that for the sort *Elem* by the following operations:

- a constant '*empty*' of sort *Container*, and
- a binary operation '*insert*' with argument sorts *Elem* and *Container*, and result sort *Container*.

This constraint would ensure that the only elements of the *Container* carrier are those obtained by a finite number of successive applications of the *insert* operation to elements of sort *Elem*, starting with the *empty* value of sort *Container*.

2.2 Structured Specifications

Structured specifications are formed from basic specifications, *references* to *named specifications*, and *instantiations* of *generic specifications*, using various constructs for composing specifications.

The semantics of a structured specification is simply a signature and a class of models.

The semantics of a structured specification is of the same form as that of a basic specification: a signature, together with a class of models. Thus the structure of a specification is *not* reflected in its models: it is used only to present the specification in a modular style. (Specification of the *architecture* of implementations is addressed in Sect. 2.3.) The symbols in the signature are called the *exported* symbols of the specification.

The interpretation of structured specification constructs involves mappings between signatures Σ , called *signature morphisms*, and corresponding mappings between models M, called *reducts along morphisms*. In CASL, a signature morphism σ from Σ to Σ' consists of a mapping which gives:

- for each sort of Σ , a corresponding sort of Σ' , preserving any subsort relationships, and
- for each operation or predicate symbol whose profile has sorts in Σ, a corresponding symbol in Σ' whose profile has the corresponding sorts, preserving any overloading between symbols whose profiles are related by subsorting. A partial operation may be mapped to a total operation, but not vice versa.

Let M' be any Σ' -model. We can define its *reduct* along the signature morphism σ to be the Σ -model M obtained as follows: each symbol of Σ is interpreted in M in exactly the same way as the corresponding symbol in Σ' is interpreted in M'. Conversely, given M, a model M' is said to be an *expansion* of M when the reduct of M' is M.

Suppose that a specification SP has signature Σ and SP' has signature Σ' . A signature morphism σ from Σ to Σ' is said to be a *specification morphism* from SP to SP' when the reduct along σ of each model of SP' is a model of SP.

For two Σ -models M_1 , M_2 , a (weak) homomorphism from M_1 to M_2 maps the elements of the carrier sets of M_1 to the elements of the corresponding carrier sets of M_2 , preserving the embeddings between subsorts, the values (and definedness) of operations, and the holding of predicates. A homomorphism is an *isomorphism* when it has an inverse homomorphism.

A model M is *initial* in a class of Σ -models if there is a unique homomorphism from M to each model in the class. When a class of models has an initial model (which need not be the case in CASL) it is unique, up to isomorphism.

With reference to the above concepts of signature morphism, model reduct, and homomorphism, we can now proceed to explain the constructs involved with structured specifications. A translation merely renames symbols.

Translating a sort symbol requires translating the profiles of all operation and predicate symbols involving that sort; translating an operation or predicate symbol has to respect overloading between symbols whose profiles are related by subsorting. The translation of sort, operation, and predicate names in CASL determines a signature morphism σ mapping the signature Σ of a specification *SP* onto a new signature Σ' . The models of the translation specification are all those models interpreting Σ' whose reducts along σ are models of *SP*.

Hiding symbols removes parts of models.

Hiding a sort symbol implies hiding also all operation and predicate symbols whose profiles involve that sort; hiding an operation or predicate symbol, however, does not have further implications. Hiding a set of symbols that occur in the signature Σ of a specification *SP* to give a subsignature Σ' determines a signature morphism σ which simply includes Σ' in Σ . The models of the hiding specification are the reducts of the models of *SP* along σ .

For example, the operation suc might be introduced purely to facilitate the specification of the natural numbers, with sort *Nat*, constants 0 and 1, and the usual arithmetic operations. Hiding *suc* removes the interpretation of *suc* from the models of the specification³ but leaves the carrier set for *Nat* unchanged.

Union of specifications identifies common symbols.

The signature of a union of specifications SP_1 , SP_2 is simply the union of their respective signatures Σ_1 , Σ_2 . The models of the union are those models of the union signature whose reducts to Σ_1 and Σ_2 along the signature inclusions satisfy SP_1 and SP_2 respectively. Thus each symbol that the two signatures have in common has a single interpretation in any model of the union specification. This is known as the 'same name, same thing' principle.

 $^{^{3}}$ The successor of a number can of course still be obtained, using addition and 1.

Extension of specifications identifies common symbols too.

The signature of the extension of a specification SP by further specification items (declarations, axioms, and constraints) is simply the extension of the signature Σ of SP with the symbols of the new declarations. The models of the extension are those models of the extended signature which satisfy the axioms and constraints specified by the extension and whose reducts to Σ satisfy SP. If the extension redeclares a symbol of SP, there is still only one occurrence of that symbol in the signature of the extension, and hence only one interpretation of it – again the 'same name, same thing' principle.

In CASL, unions, extensions, and other kinds of structured specification can be formed from specification fragments that determine only signature *extensions*, not necessarily complete signatures.

Free specifications restrict models to being free, with initiality as a special case.

When a specification is freely extended by additional specification items, the interpretations of the additional declarations are required to satisfy the axioms, but nothing more: that is, properties that are not consequences of the axioms should not hold. In particular, the domains of definition of partial operations – and the sets of arguments for which predicates hold – are as small as possible. The carriers for the original sorts are left unchanged; any new carriers are no larger than is required to provide interpretations for the operations, without unnecessary junk elements. This restriction of the models is referred to as a *freeness constraint*. In the degenerate case where the specification being enriched is empty, the models of the free extension are just the initial models.

The difference between predicates and boolean-valued operations is particularly apparent in free specifications: with predicates, it is only required to specify when they hold, since not holding is the default; with boolean-valued operations, however, the true and false values are treated symmetrically, and it is necessary to specify both cases, since neither is the default.

Generic specifications have parameters, and have to be instantiated when referenced.

A named specification may declare some *parameters*, the union of which is extended by a *body*; it is then called *generic*. The purpose of a generic specification is to reuse the body in different contexts; hence a reference to a generic specification should *instantiate* it by providing, for each parameter, an *argument specification* together with a *fitting morphism* from the parameter to the argument specification. Fitting may also be achieved by use of named *views* between the parameter and argument specifications. The instantiation of the generic specification gives the union of the arguments, together with the translation of the generic specification by an expansion of the fitting morphism. This corresponds to a so-called push-out construction – taking into account any explicit *imports* of the generic specification.

2.3 Architectural Specifications

The semantics of an architectural specification reflects its modular structure.

The intention with architectural specifications is primarily to impose structure on implementations, expressing their *composition* from component units – and thereby also a *decomposition* of the task of developing such implementations, from requirements specifications. This is in contrast to the structured specifications considered in Sect. 2.2, where the specified models have no more structure than do those of the basic specifications considered in Sect. 2.1.

Architectural specifications involve the notions of persistent function and conservative extension.

A function F mapping Σ -models to Σ' -models, where the signature Σ' extends Σ , is said to be *persistent* when for each Σ -model M, the reduct of F(M) to a Σ -model is exactly M.

A specification extension SP' of SP is said to be *conservative* when each model of SP can be expanded to a model of SP'.

A persistent function mapping models of SP to models of SP' exists only if SP' is a conservative extension of SP.

2.4 Libraries of Specifications

The semantics of a library of specifications is a mapping from the names of the specifications to their semantics.

The specification of a library gives also a library name, and determines a version number.

Getting Started

Simple specifications may be written in CASL essentially as in many other algebraic specification languages.

The simplest kind of algebraic specification is when each specified operation is to be interpreted as an ordinary *total* mathematical function: it takes values of particular types as arguments, and always returns a well-defined value. Total functions correspond to software whose execution always terminates normally. The types of values are named by simple symbols called *sorts*.

In practice, a realistic software specification involves *partial* as well as total functions. However, it may well be formed from simpler specifications, some of which involve only total functions. This chapter explains how to express such simple specifications in CASL, illustrating various features of the language.

The simple specifications discussed in this chapter can also be expressed in many previous specification languages; it is usually straightforward to reformulate them in CASL. Readers who know other specification languages will probably recognize some familiar examples among the illustrations given in this chapter.

CASL provides also useful abbreviations.

The technique of algebraic specification by axioms is generally wellsuited to expressing properties of functions. However, when functions have commonly-occurring mathematical properties, it can be tedious to give the corresponding axioms explicitly. CASL provides some useful abbreviations for such cases. Similarly, so-called free datatype declarations allow sorts and value constructors to be specified much as in functional programming languages, using a concise and suggestive notation. CASL allows loose, generated and free specifications.

The models of a loose specification include all those where the declared functions have the specified properties, without any restrictions on the sets of values corresponding to the various sorts. In models of a generated specification, in contrast, it is required that all values can be expressed by terms formed from the specified constructors, i.e. unreachable values are prohibited. In models of free specifications, it is required that values of terms are distinct except when their equality follows from the specified axioms: the possibility of unintended coincidence between them is prohibited.

Section 3.1 below focuses on loose specifications; Sect. 3.2 discusses the use of generated specifications, and Sect. 3.3 does the same for free specifications. Loose, generated, and free specifications are often used together in CASL: each style has its advantages in particular circumstances, as explained below in connection with the illustrative examples.

3.1 Loose Specifications

CASL syntax for declarations and axioms involves familiar notation, and is mostly self-explanatory.

spec STRICT_PARTIAL_ORDER = %% Let's start with a simple example ! **sort** Elem **pred** ___ < __ : Elem × Elem %% **pred** abbreviates predicate $\forall x, y, z : Elem$ • $\neg(x < x)$ %(strict)% • $x < y \Rightarrow \neg(y < x)$ %(asymmetric)% • $x < y \land y < z \Rightarrow x < z$ %(transitive)% % { Note that there may exist x, y such that neither x < y nor y < x. }%

end

The above (basic) specification, named STRICT_PARTIAL_ORDER, introduces a sort *Elem* and a binary infix predicate symbol '<'. In the declaration of a predicate symbol, argument sorts are separated by the sign ' \times ', which can be input directly as such in ISO Latin-1 or as '*' in ASCII. Note that CASL allows so-called *mixfix notation*, i.e., the specifier is free to indicate, using '__' (pairs of underscores) as *place-holders*, how to place arguments when building terms (single underscores are treated as letters in identifiers).¹ Using mixfix notation generally allows the use of familiar mathematical and programming notations, which contributes substantially to the readability of specifications.

The interpretation of the binary predicate symbol '<' is then constrained by three axioms. A set of axioms is generally presented as a 'bulleted' list of formulas, preceded by the universally quantified declaration of the relevant variables, together with their respective sorts, as shown in the above example. In CASL, axioms are written in first-order logic with equality, using quantifiers and the usual logical connectives. The universal quantification preceding a list of axioms applies to the entire list. Axioms can be annotated by labels written %(...)%, which is convenient for proper reference in explanations or by tools.

Note that ' \forall ' is input as 'forall', and that ' \bullet ' is input as '.' or ' \cdot '. The usual logical connectives ' \Rightarrow ', ' \Leftrightarrow ', ' \wedge ', ' \vee ', and ' \neg ', are input as '=>', '<=>', '/ \wedge ', ' \vee ', and 'not', respectively; ' \neg ' can also be input directly as an ISO Latin-1 character. The existential quantifier ' \exists ' is input as 'exists', and ' \exists !' is input as 'exists!'.

It is advisable to comment as appropriate the various elements introduced in a specification. The syntax for end-of-line and grouped multi-line comments is illustrated in the above example. The 'end' keyword ending a specification is optional.

The above STRICT_PARTIAL_ORDER specification is loose in the sense that it has many (non-isomorphic) models, among which models where '<' is interpreted by a total ordering relation and models where it is interpreted by a partial one.

Specifications can easily be extended by new declarations and axioms.

```
spec TOTAL_ORDER =
STRICT_PARTIAL_ORDER
then \forall x, y : Elem \bullet x < y \lor y < x \lor x = y %(total)%
end
```

Extensions, introduced by the keyword 'then', may specify new symbols, possibly constrained by some axioms, or merely require further properties of old ones, as in the above TOTAL_ORDER example, or more generally do both at the same time. In TOTAL_ORDER, we further constrain the interpretation of the predicate symbol '<' by requiring it to be a total ordering relation.

All symbols introduced in a specification are by default exported by it and visible in its extensions. This is for instance the case here for the sort *Elem* and

¹ Mixfix notation is so-called because it generalizes infix, prefix, and postfix notation to allow arbitrary mixing of argument positions and identifier tokens.

the predicate symbol '<', which are introduced in STRICT_PARTIAL_ORDER, exported by it, and therefore available in TOTAL_ORDER.²

In simple cases, an operation (or a predicate) symbol may be declared and its intended interpretation defined at the same time.

spec TOTAL_ORDER_WITH_MINMAX = TOTAL_ORDER **then ops** min(x, y : Elem) : Elem = x when x < y else y; max(x, y : Elem) : Elem = y when min(x, y) = x else x

end

TOTAL_ORDER_WITH_MINMAX extends TOTAL_ORDER by introducing two binary operation symbols *min* and *max*, for which a functional notation is to be used, so no place-holders are given. The intended interpretation of the symbol *min* is defined simultaneously with its declaration, and the same is done for *max*. For instance:

op min(x, y : Elem) : Elem = x when x < y else y

abbreviates:

op $min : Elem \times Elem \rightarrow Elem$ $\forall x, y : Elem \bullet min(x, y) = x when x < y else y$

(and similarly for *max*). As for predicate symbol declarations, in an operation symbol declaration, the argument sorts are separated by the sign ' \times '; the result sort is preceded by ' \rightarrow ', which is input as '->'.

The '... when ... else ...' construct used above is itself an abbreviation, and:

min(x, y) = x when x < y else y

abbreviates:

$$(x < y \Rightarrow \min(x, y) = x) \land (\neg(x < y) \Rightarrow \min(x, y) = y)$$

In CASL specifications, visibility is *linear*, i.e., any symbol must be declared before being used. In the above example, min should be declared before being used to define max.

Linear visibility does not imply, however, that a fixed scheme is to be used when writing specifications: the specifier is free to present the required declarations and axioms in any order, as long as the linear visibility rule is respected. For instance, one may prefer to declare first all sorts and all operation

 $^{^2}$ See Chap. 6 for constructs allowing the explicit restriction of the set of symbols exported by a specification.

or predicate symbols needed, and then specify their properties by the relevant axioms. Or, in contrast, one may prefer to have each operation or predicate symbol declaration immediately followed by the axioms constraining its interpretations. Both styles are equally fine, and can even be mixed if desired. This flexibility is illustrated in the following variant of the TOTAL_ORDER_WITH_MINMAX specification, where for explanatory purposes we refrain from using the useful abbreviations explained above.

```
spec VARIANT_OF_TOTAL_ORDER_WITH_MINMAX =

TOTAL_ORDER

then vars x, y : Elem

op min : Elem \times Elem \rightarrow Elem

• x < y \Rightarrow min(x, y) = x

• \neg(x < y) \Rightarrow min(x, y) = y

op max : Elem \times Elem \rightarrow Elem

• x < y \Rightarrow max(x, y) = y

• \neg(x < y) \Rightarrow max(x, y) = x

end
```

end

Note that in order to avoid the tedious repetition of the declaration of the variables x and y for each list of axioms, we have used here a *global* variable declaration which introduces x and y for the rest of the specification. Variable declarations are of course not exported across specification extensions: the variables x and y declared in VARIANT_OF_TOTAL_ORDER_WITH_MINMAX are not visible in any of its extensions.

Symbols may be conveniently displayed as usual mathematical symbols by means of **%display** annotations.

```
%display _<=_ %LATEX _ \leq _
```

```
spec PARTIAL_ORDER =
STRICT_PARTIAL_ORDER
then pred \_\_ \leq \__(x, y : Elem) \Leftrightarrow (x < y \lor x = y)
end
```

The above example relies on a %display annotation: while, for obvious reasons, the specification text should be *input* using the ISO Latin-1 character set, it is often convenient to *display* some symbols differently, e.g., as mathematical symbols. This is the case here where the '<=' predicate symbol is more conveniently displayed as ' \leq '. Display annotations, as any other CASL annotations, are auxiliary parts of a specification, for use by tools, and do not affect the semantics of the specification.³

 $^{^3}$ Display annotations should be provided at the beginning of a library, and are explained in more detail in Chap. 9.

In the above example, we have again used the facility of simultaneously declaring and defining a symbol (here, the predicate symbol ' \leq ') in order to obtain a more concise specification.

The %*implies* annotation is used to indicate that some axioms are supposedly redundant, being consequences of others.

```
spec PARTIAL_ORDER_1 =
PARTIAL_ORDER
then %implies
\forall x, y, z : Elem \bullet x \le y \land y \le z \Rightarrow x \le z %(transitive)%
end
```

The %implies annotation above is used to emphasize that the transitivity of ' \leq ' should follow from the other axioms, or, in other words, that the model class of PARTIAL_ORDER_1 is exactly the same as the model class of PARTIAL_ORDER. The %implies annotation applies to the whole of the specification extension where it occurs (which happens here to introduce a single axiom).

Note however that an annotation does not affect the semantics of a specification, hence removing the **%implies** annotation does not change the class of models of the above specification. The sole aim of an **%implies** annotation is to stress the specifier's intentions, and it will also help readers confirm their understanding. Some tools may of course use such annotations to generate corresponding proof obligations. For instance, here, the proof obligation is:

PARTIAL_ORDER $\models \forall x, y, z : Elem \bullet x \le y \land y \le z \Rightarrow x \le z$

Discharging these proof obligations increases the trustworthiness of a specification.

To fully understand that an **%implies** annotation has no effect on the semantics, the best is to consider an example where the corresponding proof obligation cannot be discharged, as shown below.

```
spec IMPLIES_DOES_NOT_HOLD =
PARTIAL_ORDER
then %implies
\forall x, y : Elem \bullet x < y \lor y < x \lor x = y %(total)%
end
```

Since the loose specification PARTIAL_ORDER has models where '<' is interpreted by a partial ordering relation, the proof obligation corresponding to the above **%implies** annotation cannot be discharged. However, since annotations have no impact on the semantics, the specification IMPLIES_DOES_ NOT_HOLD is well-formed and just constrains the interpretation of '<' to be a total ordering relation. The fact that the proof obligation cannot be discharged merely points out a potential mistake in the specification. Attributes may be used to abbreviate axioms for associativity, commutativity, idempotence, and unit properties.

```
spec MONOID =

sort Monoid

ops 1 : Monoid;

\_*\_: Monoid × Monoid → Monoid, assoc, unit 1
```

end

The above example introduces a constant symbol 1 of sort *Monoid*, then a binary operation symbol '*', which is asserted to be associative and to have 1 as unit element. (Note that there is no ' \rightarrow ' sign before the sort when declaring a constant.) The *assoc* attribute abbreviates, as expected, the following axiom:

$$\forall x, y, z : Monoid \bullet (x * y) * z = x * (y * z)$$

The 'unit 1' attribute abbreviates:

 $\forall x: Monoid \bullet (x * 1 = x) \land (1 * x = x)$

Note that to make the use of 'unit 1' legal, it is necessary to have previously declared the constant 1, to respect the linear visibility rule.

Other available attributes are *comm*, which abbreviates the obvious axiom stating that a binary operation is commutative, and *idem*, which can be used to assert the idempotence of a binary operation f (i.e., that f(x, x) = x).

Asserting '*' to be associative using the attribute *assoc* makes the term x * y * z well-formed (assuming x, y, z of the right sort), while otherwise grouping parentheses would be required. Moreover, it is expected that some tools (e.g., systems based on rewriting) may make special use of the *assoc* attribute, so it is generally advisable to use this attribute instead of stating the same property by an axiom (the same applies to the other attributes).

Genericity of specifications can be made explicit using parameters.

```
spec GENERIC_MONOID [sort Elem] =

sort Monoid

ops inj : Elem \rightarrow Monoid;

1 : Monoid;

\_*\_: Monoid \times Monoid \rightarrow Monoid, assoc, unit 1

\forall x, y : Elem \bullet inj(x) = inj(y) \Rightarrow x = y
```

```
end
```
The above example describes monoids built over arbitrary elements (of sort *Elem*). The intention here is to reuse the specification GENERIC_MONOID to derive from it specifications of monoids built over, say, characters, symbols, etc. In such cases, it is appropriate to emphasize the intended genericity of the specification by making explicit, in a distinguished *parameter part* (which is here [**sort** *Elem*]), the piece of specification that is intended to vary in the derived specifications. In these, it will then be possible to *instantiate* the parameter part as desired in order to specialize the specification as appropriate (to obtain, e.g., a specification of monoids built over characters). A named specification with one or more parameter(s) is called *generic*.

The body of the generic specification GENERIC_MONOID is an extension of what is specified in the parameter part. Hence an alternative to the above generic specification GENERIC_MONOID is the following, less elegant, nongeneric specification (which cannot be specialized by instantiation):

```
spec NON_GENERIC_MONOID =

sort Elem

then sort Monoid

ops inj : Elem \rightarrow Monoid;

1 : Monoid;

\_*\_: Monoid \times Monoid \rightarrow Monoid, assoc, unit 1

\forall x, y : Elem \bullet inj(x) = inj(y) \Rightarrow x = y
```

end

A generic specification may have more than one parameter, and parameters can be arbitrary specifications, named or not. When reused by reference to its name, a generic specification must be instantiated. Generic specifications and how to instantiate them are discussed in detail later in Chap. 7. Using generic specifications when appropriate improves the reusability of specification definitions.

 $References\ to\ generic\ specifications\ always\ instantiate\ the\ parameters.$

```
spec GENERIC_COMMUTATIVE_MONOID [sort Elem] = GENERIC_MONOID [sort Elem]
then \forall x, y : Monoid \bullet x * y = y * x
end
```

The above (generic) specification GENERIC_COMMUTATIVE_MONOID is defined as an extension of GENERIC_MONOID, which should therefore be instantiated, as explained above. Instantiating a generic specification is done by providing an argument specification that 'fits' the parameter part of the generic specification to be instantiated. It is however quite frequent that the instantiation is 'trivial', i.e., the argument specification is identical to the parameter one. This is the case for the above example, where the generic specification GENERIC_MONOID is instantiated by providing the same argument specification '**sort** *Elem*' as the original parameter.

```
spec GENERIC_COMMUTATIVE_MONOID_1 [sort Elem] = GENERIC_MONOID [sort Elem]
then op \__*\_: Monoid \times Monoid \rightarrow Monoid, comm end
```

GENERIC_COMMUTATIVE_MONOID_1 is an alternative version of the former specification where, instead of requiring explicitly with an axiom the commutativity of the operation '*', we require it using the attribute *comm*. As explained before, it is in general better to describe such requirements using attributes rather than explicit axioms, since it is expected that some tools will rely on these attributes for specialized algorithms (e.g., AC term rewriting).

This example illustrates also an important feature of CASL, the 'same name, same thing' principle. The operation symbol '*' is indeed declared twice, with the same profile, first in GENERIC_MONOID and then again in GENERIC_COMMUTATIVE_MONOID_1 (the second declaration being enriched by the attribute comm). This is perfectly fine and defines only one binary operation symbol '*' with the corresponding profile, according to the 'same name, same thing' principle. This principle applies to sorts, as well as to operation and predicate symbols. It applies both to symbols defined locally and to symbols imported from an extended specification, as it is the case here for '*'. Of course, it does not apply between separate named specifications, i.e., the same symbol may be used in different named specifications with entirely different interpretations.

Note that for operation and predicate symbols, the 'same name, same thing' principle is a little more subtle than for sorts: the 'name' of an operation (or of a predicate) includes its profile of argument and result sorts, so two operations defined with the same symbol but with different profiles do not have the same 'name', the symbol is just *overloaded*. When an overloaded symbol is used, the intended profile is to be determined by the context (e.g., the sorts of the arguments to which the symbol is applied).⁴ Explicit disambiguation can be used when needed, by specifying the profile (or result sort) in an application.⁵ Note that overloaded constants are allowed in CASL (e.g., *empty* may be declared to be a constant of various sorts of collections).

 $^{^4}$ See also the discussion of overloading in presence of subsorts in Chap. 5, p. 61.

⁵ For instance, depending on the context, the term t1 * t2 can be disambiguated by writing $(op * : Monoid \times Monoid \rightarrow Monoid)(t1, t2)$, or just (t1 : Monoid)*(t2 : Monoid), or even (t1 * t2) : Monoid.

Datatype declarations may be used to abbreviate declarations of sorts and constructors.

spec CONTAINER [**sort** Elem] = **type** Container ::= empty | insert(Elem; Container) **pred** __is_in__: $Elem \times Container$ $\forall e, e' : Elem; C : Container$ • $\neg(e \ is_in \ empty)$ • $e \ is_in \ insert(e', C) \Leftrightarrow (e = e' \lor e \ is_in \ C)$

end

Specifications of 'datatypes' with constructors are frequently needed. CASL provides special constructs for datatype declarations to abbreviate the corresponding rather tedious declarations. For instance, the above datatype declaration:

type Container ::= empty | insert(Elem; Container)

abbreviates:

A datatype declaration looks like a context-free grammar in a variant of BNF. It declares the symbols on the left of '::=' as sorts, and for each alternative on the right it declares a constructor.

A datatype declaration as the one above is *loose* since it does not imply any constraint on the values of the declared sorts: there may be some values of sort *Container* that are not built by any of the declared constructors, and the same value may be built by different applications of the constructors to some arguments.

Datatype declarations may also be specified as generated (see Sect. 3.2) or free (see Sect. 3.3). Moreover, selectors, which are usually partial operations, may be specified for each component (see Chap. 4).

Loose datatype declarations are appropriate when further constructors may be added in extensions.

```
spec MARKING_CONTAINER [sort Elem] =
    CONTAINER [sort Elem]
then type Container ::= mark_insert(Elem; Container)
```

pred __is_marked_in_ : $Elem \times Container$ $\forall e, e' : Elem; C : Container$

- $e \text{ is_in mark_insert}(e', C) \Leftrightarrow (e = e' \lor e \text{ is_in } C)$
- $\neg(e \ is_marked_in \ empty)$
- $e \ is_marked_in \ insert(e', C) \Leftrightarrow e \ is_marked_in \ C$
- $e \text{ is_marked_in mark_insert}(e', C) \Leftrightarrow (e = e' \lor e \text{ is_marked_in } C)$

end

The above specification extends CONTAINER (trivially instantiated) by introducing another constructor *mark_insert* for the sort *Container* (hence, values added to a container may now be 'marked' or not). Note that we heavily rely on the 'same name, same thing' principle here, since it ensures that the sort *Container* introduced by the datatype declaration of CONTAINER and the sort *Container* introduced by the datatype declaration of MARKING_ CONTAINER are the same sort, which implies that the combination of both datatype declarations is equivalent to:

Note that since 'new' values may be constructed by *mark_insert*, it is necessary to extend the specification of the predicate symbol *is_in* by an extra axiom taking care of the newly introduced constructor.

3.2 Generated Specifications

Sorts may be specified as generated by their constructors.

```
spec GENERATED_CONTAINER [sort Elem] =

generated type Container ::= empty | insert(Elem; Container)

pred __is_in__: Elem × Container

\forall e, e' : Elem; C : Container

• \neg(e \ is\_in \ empty)

• e \ is\_in \ insert(e', C) \Leftrightarrow (e = e' \lor e \ is\_in \ C)
```

end

When a datatype is declared as *generated*, as in the above example, the corresponding sort is constrained to be generated by the declared constructors, which means that any value of this sort is built by application of constructors. This constraint is sometimes referred to as the 'no junk' principle. For instance, in the above example, having declared the datatype *Container* to be generated entails that in any model of GENERATED_CONTAINER, any value of

sort *Container* is denotable by a term built with *empty*, *insert*, and variables of sort *Elem* only.

As a consequence, properties of values of sort *Container* can be proved by induction on the declared constructors. A major benefit of generated datatypes is indeed that induction on the declared constructors is a sound proof principle.

The construct 'generated type ...' used above is just an abbreviation for 'generated { type ... }', and 'generated' can be used around arbitrary signature declarations, enclosed in '{' and '}'.

Generated specifications are in general loose.

spec	Generated_Container_Merge [sort <i>Elem</i>] =				
	GENERATED_CONTAINER [sort Elem]				
\mathbf{then}	op $_merge_: Container \times Container \rightarrow Container$				
	$\forall e : Elem; C, C' : Container$				
	• $e \text{ is_in } (C \text{ merge } C') \Leftrightarrow (e \text{ is_in } C \lor e \text{ is_in } C')$				

end

A generated specification is in general loose. For instance, GENERATED_ CONTAINER is loose since, although all values of sort *Container* are specified to be generated by *empty* and *insert*, the behavior of the *insert* constructor is still loosely specified (nothing is said about the case where an element is inserted into a container which already contains this element). Hence GENERATED_ CONTAINER admits several non-isomorphic models.

GENERATED_CONTAINER_MERGE is as loose as GENERATED_CONTAINER with respect to *insert*, and in addition, the newly introduced operation symbol *merge* is also loosely specified: nothing is said about what happens when merging two containers which share some elements.

It is important to understand that looseness of a specification is not a problem, but on the contrary avoids unnecessary overspecification. In particular, loose specifications are in general well-suited to capturing requirements.

The fact that *merge* is loosely specified does not mean that it can produce new values of the sort *Container*. On the contrary, since this sort has been specified as being generated by *empty* and *insert*, it follows that any value denotable by a term of the form $merge(\ldots,\ldots)$ can also be denoted by a term built with *empty* and *insert* (and no *merge*). Hence, for the specification GENERATED_CONTAINER_MERGE, proofs by induction on *Container* only need to consider *empty* and *insert*, and not *merge*, as was the case for GENERATED_CONTAINER.

Generated specifications need not be loose.

spec GENERATED_SET [**sort** Elem] = generated type $Set ::= empty \mid insert(Elem; Set)$ pred __is_in__: $Elem \times Set$ **ops** $\{_\}(e: Elem): Set = insert(e, empty);$: $Set \times Set \rightarrow Set;$ __ U __ $: Elem \times Set \rightarrow Set$ remove $\forall e, e' : Elem; S, S' : Set$ • $\neg(e \ is_in \ empty)$ • $e \text{ is_in insert}(e', S) \Leftrightarrow (e = e' \lor e \text{ is_in } S)$ • $S = S' \Leftrightarrow (\forall x : Elem \bullet x \text{ is_in } S \Leftrightarrow x \text{ is_in } S')$ %(equal_sets)% • $e \text{ is_in } (S \cup S') \Leftrightarrow (e \text{ is_in } S \lor e \text{ is_in } S')$ • $e \text{ is_in remove}(e', S) \Leftrightarrow (\neg(e = e') \land e \text{ is_in } S)$ then %implies $\forall e, e' : Elem; S : Set$ • insert(e, insert(e, S)) = insert(e, S)• insert(e, insert(e', S)) = insert(e', insert(e, S))generated type $Set ::= empty | \{_\}(Elem) | _\cup _(Set; Set) |$ **op** $__\cup_:$ Set \times Set \rightarrow Set, assoc, comm, idem, unit empty end

Although generated specifications are in general loose, they need not be so, as illustrated by the above GENERATED_SET specification, where the axiom %(equal_sets)%, combined with the axioms defining *is_in*, fully constraints (up to isomorphism) the interpretations of the sort Set and of the constructors *empty* and *insert*, once an interpretation for the sort *Elem* is chosen.

Note also that this example displays the power of the annotation **%implies**. Remember that this annotation applies to the whole of the specification extension where it occurs, so here it applies not only to the two explicit axioms about *insert*, but also to the properties corresponding to the attributes of \cup as well as to the generatedness constraint. Hence, the %implies annotation is used here not only to stress that the usual properties of *insert* are expected to follow from the preceding declarations and axioms, but also that an alternative induction scheme, based on *empty*, $\{ _ \}$ and $_ \cup _$, can be used for sets. Moreover, it asserts that $_ \cup _$ is expected to be associative, commutative, idempotent (i.e., $S \cup S = S$), and to have *empty* as unit. Note again that this % implies part heavily relies on the 'same name, same meaning' principle.

Generated types may need to be declared together.

The following specification fragment illustrates what may go wrong.

sort Node
generated type Tree ::= mktree(Node; Forest)
generated type Forest ::= empty | add(Tree; Forest)

The above is *incorrect*, due to the linear visibility rule. This can easily be fixed by replacing '**sort** Node' by '**sorts** Node, Tree, Forest'. Even when corrected, the above is *wrong*, since the corresponding semantics is not what a naive reader may expect. One may expect that only models where the carrier sets of the sorts Tree and Forest are generated by mktree, empty and add are acceptable, but more models satisfy the above two separate sort generatedness constraints. For instance, a model with both a junk tree *jt* and a junk forest *jf* fulfills the above declarations (assuming that the interpretations of mktree and add on *jt* and *jf* in this model are such that jt = mktree(n, jf) for any node *n* and that jf = add(jt, jf)). Hence, one must write instead:

sort Node generated types Tree ::= mktree(Node; Forest); $Forest ::= empty \mid add(Tree; Forest)$

Here, the mutually recursive datatypes *Tree* and *Forest* are correctly defined simultaneously within the same **generated types** construct, and the resulting semantics is the expected one (without junk values for trees and forests). Note that therefore, the linear visibility rule is *not* applicable *within* a **generated types** construct (to allow such mutually recursive definitions), but that this is the only exception to the linear visibility principle. Only mutually recursive generated datatypes need to be declared together; in simpler cases, it makes no difference to have a sequence of successive generated datatypes.⁶

3.3 Free Specifications

Free specifications provide initial semantics and avoid the need for explicit negation.

spec NATURAL = **free type** $Nat ::= 0 \mid suc(Nat)$

⁶ The same explanations apply to free datatypes, introduced in the next subsection.

A free datatype declaration corresponds to the so-called 'no junk, no confusion' principle: there are no other values of sort Nat than those denoted by the constructor terms built with θ and suc, and all distinct constructor terms denote different values.

Hence, a **free** datatype declaration such as the one above has *exactly* the same effect as the corresponding **generated** datatype declaration, *together* with axioms stating that *suc* is injective, and that θ cannot be the successor of a natural number. An alternative to the above 'free type $Nat ::= \theta \mid suc(Nat)$ ' is therefore:

generated type $Nat ::= 0 \mid suc(Nat)$ $\forall x, y : Nat \bullet suc(x) = suc(y) \Rightarrow x = y$ $\forall x : Nat \bullet \neg(0 = suc(x))$

Free datatype declarations are particularly convenient for defining enumerated datatypes.

```
spec COLOR =
    free type RGB ::= Red | Green | Blue
    free type CMYK ::= Cyan | Magenta | Yellow | Black
end
```

end

Using 'free' instead of 'generated' for defining enumerated datatypes saves the writing of many explicit distinctness assertions (for instance, here, $\neg(Red = Green), \neg(Red = Blue), \ldots$).

Free specifications can also be used when the constructors are related by some axioms.

```
spec INTEGER =

free { type Int ::= 0 | suc(Int) | pre(Int)

\forall x : Int \bullet suc(pre(x)) = x

\bullet pre(suc(x)) = x }
```

end

When some relations are to be imposed between the constructors (as is the case here for *suc* and *pre* which are inverses of each other), a **free** datatype declaration cannot be used, since the contradiction between the 'no confusion' principle and the axioms imposed on the constructors would lead to an inconsistent specification. Instead, one should impose a '*freeness constraint*' around the datatype declaration followed by the required axioms. A freeness

constraint, expressed by the keyword $\mathbf{free},$ can be imposed around any specification.

In the case of the above INTEGER specification, the freeness constraint imposes that the semantics of the specification is the class of all algebras isomorphic to the quotient of the constructor terms by (the minimal congruence induced by) the given axioms. This is exactly the desired semantics. More generally, a freeness constraint around a specification indicates its initial model, which may not exist, of course. It is however well-known that initial models of basic specifications with axioms restricted to Horn clauses (of which equations as in INTEGER are a special case) always exist.⁷ Remember also that equality holds minimally in initial models of equational specifications.

Predicates hold minimally in models of free specifications.

```
spec NATURAL_ORDER =
NATURAL
then free { pred _ < _ : Nat \times Nat
\forall x, y : Nat
• 0 < suc(x)
• x < y \Rightarrow suc(x) < suc(y) }
```

\mathbf{end}

A freeness constraint imposed around a predicate declaration followed by some defining axioms has the effect that the predicate only holds when this follows from the given axioms, and does not hold otherwise. For instance, in the above example, it is not necessary to explicitly state that ' $\neg(0 < 0)$ ', since this will follow from the imposed freeness constraint. Hence, in such cases a freeness constraint has exactly the same effect as the so-called 'negation as failure' or 'closed world assumption' principles in logic programming. More generally, it is often convenient to define a predicate within a freeness constraint, since by doing so, one has to specify the 'positive' cases only.

Operations and predicates may be safely defined by induction on the constructors of a free datatype declaration.

spec Natural_Arithmetic = Natural_Order

⁷ Strictly speaking, existence of initial models depends on a further requirement: namely the existence of a ground term for each sort. This ensures that the term-algebra has non-empty carriers and hence is a CASL model.

then ops 1 : $Nat = suc(\theta)$; $_+_: Nat \times Nat \rightarrow Nat, assoc, comm, unit 0;$ $_*_: Nat \times Nat \rightarrow Nat, assoc, comm, unit 1$ $\forall x, y : Nat$ • x + suc(y) = suc(x + y)• $x * \theta = \theta$ • x * suc(y) = (x * y) + xend

To define some operation on a free datatype, it is generally recommended to make a case distinction with respect to the various constructors defined. This is illustrated by the above definitions of + and + (although for the '+' operation, the case for the constructor θ is already taken care of by the attribute 'unit 0').⁸

More care may be needed when defining operations or predicates on free datatypes when there are axioms relating the constructors.

```
spec INTEGER_ARITHMETIC =
       INTEGER
then ops 1 : Int = suc(\theta);
            \_+\_: Int \times Int \rightarrow Int, assoc, comm, unit 0;
            \_- \_: Int \times Int \rightarrow Int;
            \_*\_: Int \times Int \rightarrow Int, assoc, comm, unit 1
       \forall x, y : Int
        • x + suc(y) = suc(x + y)
       • x + pre(y) = pre(x + y)
        • x - 0 = x
        • x - suc(y) = pre(x - y)
        • x - pre(y) = suc(x - y)
        • x * \theta = \theta
       • x * suc(y) = (x * y) + x
       • x * pre(y) = (x * y) - x
```

end

While a case distinction with respect to the constructors of a free datatype is harmless, this may not be the case for a datatype defined within a freeness constraint, since, due to the axioms relating the constructors to each other, some cases may overlap. This does not mean, however, that one cannot use the case distinction, but just that more attention should be paid than for a free datatype – since one needs to ensure that the definitions lead to the same

⁸ In specification libraries, ordinary decimal notation for natural numbers can be provided by use of so-called literal syntax annotations, see Chap. 9.

results for overlapping cases. For instance, in the above example no problem arises. But one should be more careful with the next one, since a negative integer can be of the form suc(x), hence asserting, e.g., $\theta \leq suc(x)$, would of course be wrong.

```
spec INTEGER_ARITHMETIC_ORDER =
INTEGER_ARITHMETIC
then preds ______, ______, ______, ______. : Int × Int
\forall x, y : Int
• 0 \le 0
• \neg (0 \le pre(0))
• 0 \le x \Rightarrow 0 \le suc(x)
• \neg (0 \le x) \Rightarrow \neg (0 \le pre(x))
• suc(x) \le y \Leftrightarrow x \le pre(y)
• pre(x) \le y \Leftrightarrow x \le suc(y)
• x \ge y \Leftrightarrow y \le x
• x < y \Leftrightarrow (x \le y \land \neg (x = y))
• x > y \Leftrightarrow y < x
```

end

Generic specifications often involve free extensions of (loose) parameters.

spec LIST [**sort** *Elem*] = **free type** *List* ::= *empty* | *cons*(*Elem*; *List*)

The parameter of a generic specification should be loose to cope with the various expected instantiations. On the other hand, it is a frequent situation that the body of the generic specification should have a free, initial interpretation. This is illustrated by the above example, where we want to combine a loose interpretation for the sort *Elem* with a free interpretation for lists. The following example is similar in spirit.

```
spec SET [sort Elem] =
    free { type Set ::= empty | insert(Elem; Set)
        pred __is_in__ : Elem × Set
        \forall e, e' : Elem; S : Set
        insert(e, insert(e, S)) = insert(e, S)
        insert(e, insert(e', S)) = insert(e', insert(e, S))
        \neg(e \ is_in \ empty)
        e \ is_in \ insert(e, S)
        e \ is_in \ insert(e', S) \ if \ e \ is_in \ S }
```

end

As for the LIST example, we want to have a loose interpretation for the sort *Elem* and a free interpretation for sets. Since some axioms are required to hold for the *Set* constructors *empty* and *insert*, we cannot use a free datatype declaration, hence we use a freeness constraint.

Note that since, as already explained, predicates hold minimally in models of free specifications, it would have been enough, in the above example, to define the predicate is_in by the sole axiom e is_in insert(e, S).⁹ However, doing so would have decreased the comprehensibility of the specification and this is the reason why we have preferred a more verbose axiomatization of the predicate is_in .

Note also the use of the keyword `if" to write an implication in the reverse order:

 $e \text{ is_in insert}(e', S) \text{ if } e \text{ is_in } S$

is equivalent to:

 $e \text{ is_in } S \Rightarrow e \text{ is_in } insert(e', S)$

The following example specifies the transitive closure of an arbitrary binary relation R on some sort *Elem* (both provided by the parameter).

spec TRANSITIVE_CLOSURE [**sort** Elem **pred** __R_ : Elem × Elem] = **free** { **pred** __R⁺_ : Elem × Elem $\forall x, y, z : Elem$ • $x \ R \ y \Rightarrow x \ R^+ y$ • $x \ R^+ y \land y \ R^+ z \Rightarrow x \ R^+ z$ }

In the above example, it is crucial that predicates hold minimally in models of free specifications, since this property ensures that what we define as ' R^+ ' is actually the *smallest* transitive relation including R. Without requiring the freeness constraint, one would allow arbitrary transitive relations containing R (and these undesired relations cannot be eliminated merely by specifying further first-order axioms).

Loose extensions of free specifications can avoid overspecification.

spec NATURAL_WITH_BOUND =
 NATURAL_ARITHMETIC
then op max_size : Nat

• $0 < max_size$

end

⁹ If an element e belongs to a set S', then this set S' can always be denoted by a constructor term of the form insert(e, S), due to the axioms constraining the constructor *insert*.

The above example shows another benefit of mixing loose and initial semantics. Assume that at this stage we want to introduce some bound, of sort Nat, without fixing its value vet (this value is likely to be fixed later in some refinement, and all that we need for now is the existence of some bound). This is provided by the above specification NATURAL_WITH_BOUND, where we mix an initial interpretation for the sort *Nat* (defined using a free datatype declaration in NATURAL) and a loose interpretation for the constant max_size. Each model of NATURAL_WITH_BOUND will provide a fixed interpretation of the constant max_size, and all these models are captured by NATURAL_ WITH_BOUND, which is in this sense loose. Using such loose extensions is in general appropriate to avoid unnecessary overspecification.

spec Set_CHOOSE [**sort** *Elem*] = SET [sort Elem] then op $choose : Set \rightarrow Elem$ $\forall S: Set \bullet \neg (S = empty) \Rightarrow choose(S) is_in S$ end

This example shows again the benefit of mixing initial and loose semantics. Here, we want to extend sets, defined using a free constraint in SET, by a loosely specified operation *choose*.¹⁰ At this stage, the only property required for *choose* is to provide some element belonging to the set to which it is applied, and we do not want to specify more precisely which specific element is to be chosen. Note that each model of SET_CHOOSE will provide a function implementing some specific choice strategy, and that since all these interpretations of *choose* have to be *functions*, they are necessarily 'deterministic' (e.g., applied twice to the same set argument, they return the same result).

Datatypes with observer operations or predicates can be specified as generated instead of free.

spec Set_GENERATED [**sort** Elem] = generated type $Set ::= empty \mid insert(Elem; Set)$ **pred** $_is_in_: Elem \times Set$ $\forall e, e' : Elem; S, S' : Set$ • $\neg(e \ is_in \ empty)$ • $e \text{ is_in insert}(e', S) \Leftrightarrow (e = e' \lor e \text{ is_in } S)$ • $S = S' \Leftrightarrow (\forall x : Elem \bullet x \text{ is_in } S \Leftrightarrow x \text{ is_in } S')$

end

 $^{^{10}}$ For the purpose of this example, we disregard the fact that choose should be undefined on the empty set, and we just leave this case unspecified. Partial functions are discussed in Chap. 4.

The above specification is an alternative to the specification SET (see p. 40). Both SET and SET_GENERATED define exactly the same class of models. The former specification relies on a freeness constraint, while SET_GENERATED relies on the observer is_in to specify when two sets are equal. Indeed, the last axiom of SET_GENERATED expresses directly that two sets having exactly the same elements are equal values. This axiom, together with the first two axioms defining is_in , will entail as well the expected properties on the constructor *insert* (see GENERATED_SET p. 35). Note also that since, in SET_GENERATED, the predicate is_in is not defined within a freeness constraint, we specify when it holds using ' \Leftrightarrow ' rather than a one-way implication.

While a freeness constraint may be unavoidable to define a predicate, as illustrated by TRANSITIVE_CLOSURE, the choice between relying on a freeness constraint to define a datatype such as *Set*, or using instead a generated datatype declaration together with some observers to unambiguously determine the values of interest, is largely a matter of convenience. One may argue that SET is more suitable for prototyping tools based on term rewriting, while SET_GENERATED is more suitable for theorem-proving tools.

The %def annotation is useful to indicate that some operations or predicates are uniquely defined.

```
spec SET_UNION [sort Elem ] =

SET [sort Elem]

then %def

ops \_ \cup \_ : Set × Set \rightarrow Set, assoc, comm, idem, unit empty;

remove : Elem × Set \rightarrow Set

\forall e, e' : Elem; S, S' : Set

• S \cup insert(e', S') = insert(e', S \cup S')

• remove(e, empty) = empty

• remove(e, insert(e, S)) = remove(e, S)

• remove(e, insert(e', S)) = insert(e', remove(e, S)) if \neg(e = e')
```

end

The annotation **%def** expresses that SET_UNION is a *definitional extension* of SET, i.e., that each model of SET can be *uniquely* extended to a model of SET_UNION, which means that the operations introduced in SET_UNION are uniquely defined. As with the **%implies** annotation, the **%def** annotation has no impact on the semantics, but a corresponding proof obligation can be generated, to be discharged by theorem proving tools. The **%def** annotation is especially useful to stress that the specifier's intention is to impose a unique interpretation of what is defined within the scope of this annotation (once an interpretation for the part which is extended has been chosen).

Operations can be defined by axioms involving observer operations, instead of inductively on constructors.

```
spec Set_UNION_1 [sort Elem] =
         SET_GENERATED [sort Elem]
then %def
         ops \_ \cup \_ : Set \times Set \rightarrow Set, assoc, comm, idem, unit empty;
                remove: Elem \times Set \rightarrow Set
         \forall e, e' : Elem; S, S' : Set
          • e \text{ is_in } (S \cup S') \Leftrightarrow (e \text{ is_in } S \lor e \text{ is_in } S')
          • e \text{ is_in remove}(e', S) \Leftrightarrow (\neg(e = e') \land e \text{ is_in } S)
```

end

The specification SET_UNION_1 is an alternative to SET_UNION and defines exactly the same model class. While an inductive definition style was chosen for the operations ' \cup ' and *remove* in SET_UNION, in SET_UNION_1 these operations are defined 'implicitly' by characterizing their results through the observer *is_in*. Note that this 'observer' style does not prevent us providing a unique definition of both operations, as claimed by the %def annotation.

Similarly to the discussion on the respective merits of SET and of SET_ GENERATED, the choice between an inductive definition style and an 'observer' definition style is partly a matter of taste. One may argue that the 'observer' definition style is more abstract in the sense that there is no hint to any algorithmic computation of the so-defined operations, while the inductive definition style mimics a recursive definition in a functional programming language. Again, the inductive definition style may be more suitable for prototyping tools based on term rewriting, while the 'observer' definition style may be more suitable for theorem-proving tools.

Sorts declared in free specifications are not necessarily generated by their constructors.

```
spec UNNATURAL =
       free { type UnNat ::= 0 | suc(UnNat)
                op \_-+\_: UnNat \times UnNat \rightarrow UnNat,
                                assoc, comm, unit 0
                \forall x, y : UnNat \bullet x + suc(y) = suc(x + y)
                \forall x : UnNat \bullet \exists y : UnNat \bullet x + y = 0 \}
```

end

This rather peculiar example illustrates the fact that a sort defined within a freeness constraint need not be generated by its constructors. In UNNAT-URAL, the specification enclosed within the free $\{ \dots \}$ construct specifies Abelian groups with one generator $suc(\theta)$, and the integers are the free such Abelian group. Hence, the (unique up to isomorphism) model of UNNATURAL corresponds to integers, and not to natural numbers as one may expect – just consider the last axiom. This example points out why in general datatypes defined using freeness constraints can be more difficult to understand than datatypes defined using generatedness constraints. However, the reader should be aware that the specification UNNATURAL uses a proper first-order formula with an existential quantifier in the axioms. The specification UNNATURAL is provided here for explanatory purposes only, and clearly the writing of similar specifications should be discouraged. When only Horn clauses are used as axioms in a freeness constraint, then the datatype will indeed be generated by its constructors.

Partial Functions

Partial functions arise naturally.

Partial functions arise in a number of situations. CASL provides means for the declaration of partial functions, the specification of their domains of definition, and more generally the specification of system properties involving partial functions. The aim of this chapter is to discuss and illustrate how to handle partial functions in CASL specifications.

4.1 Declaring Partial Functions

Partial functions are declared differently from total functions.

```
spec SET_PARTIAL_CHOOSE [sort Elem] = GENERATED_SET [sort Elem]
then op choose : Set \rightarrow? Elem
end
```

The *choose* function on sets is naturally a partial function, expected to be undefined on the empty set. In CASL, a partial function is declared similarly to a total one, *but* for the question mark '?' following the arrow in the profile. It is therefore quite easy to distinguish the functions declared to be total from the ones declared to be partial.

A function declared to be partial may happen to be total in some of the models of the specification. For instance, the above specification SET_ PARTIAL_CHOOSE does not exclude models where the function symbol *choose* is interpreted by a total function, defined on all set values. Axioms can be

[©] IFIP International Federation for Information Processing 2004

used to specify the domain of definition of a partial function, and how to do this is detailed later in this chapter.

Terms containing partial functions may be undefined, i.e., they may fail to denote any value.

For instance, the (value of the) term choose(empty) may be undefined.¹ This is more natural than insisting that choose(empty) has to denote some arbitrary but fixed element of *Elem*.

Note that variables range only over defined values, and therefore a variable always denotes a value, in contrast to terms containing partial functions.

Functions, even total ones, propagate undefinedness.

If the term choose(S) is undefined for some value of S, then the term insert(choose(S), S') is undefined as well for this value of S, although *insert* is a total function.

Predicates do not hold on undefined arguments.

CASL is based on classical two-valued logic. A predicate symbol is interpreted by a relation, and when the value of some argument term is undefined, the application of a predicate to this term does not hold. For instance, if the term choose(S) is undefined, then the atomic formula choose(S) is_in S does not hold.

Equations hold when both terms are undefined.

In CASL, equations are by default *strong*, which means that they hold not only when both sides denote equal values, but also when both sides are simultaneously undefined. For instance, let us consider the equation:

insert(choose(S), insert(choose(S), empty)) = insert(choose(S), empty)

¹ Note that the term *choose(empty)* is well-formed and therefore is a 'correct term'. It is *its value* which may be undefined. To avoid unnecessary pedantry, in the following we will simply write that a term is undefined to mean that its value is so. Obviously, a term with variables may be defined for some values of the variables and undefined for other values.

Either choose(S) is defined and then both sides are defined and denote equal values due to the axioms on *insert*, or choose(S) is undefined and then both sides are undefined, and the strong equation 'holds trivially'.

CASL provides also so-called *existential equations*, explained at the end of this chapter.

Special care is needed in specifications involving partial functions.

Partial functions are intrinsically more difficult to understand and specify than total ones. This is why special care is needed when writing the axioms of specifications involving partial functions. The point is that an axiom may imply the definedness of terms containing partial functions, and as a consequence that these functions are total, which may not be what the specifier intended. Here are three typical cases:

- Asserting choose(S) is_in S as an axiom implies that choose(S) is defined, for any S. The point here is that since predicates applied to an undefined term do not hold, in any model satisfying choose(S) is_in S, the function choose must be total (i.e., always defined).
- Asserting remove(choose(S), insert(choose(S), empty)) = empty as an axiom implies that choose(S) is defined for any S, since the term empty is always defined. To understand this, assume that choose is undefined for some set value of S; then the above equation cannot hold for this value, since the undefinedness of choose(S) implies the undefinedness of remove(choose(S), insert(choose(S), empty)), giving a contradiction with the definedness of empty. Hence, an equation between a term involving a partial function PF and a term involving total functions only may imply that the partial function PF is always defined.
- Asserting insert(choose(S), S) = S as an axiom implies that choose(S) is defined for any S, since a variable always denotes a defined value. This case is indeed similar to the previous one, the only difference being that now the right-hand side of the equation is a variable (instead of a term involving total functions only).

Moreover, the 'same name, same thing' principle has a subtle side-effect regarding partial operations: if an operation is declared both as a total operation and as a partial operation with the same profile (i.e., the same argument sorts and the same result sort) then it is interpreted as a total operation in all models of the specification.

4.2 Specifying Domains of Definition

The definedness of a term can be checked or asserted.

```
spec SET_PARTIAL_CHOOSE_1 [sort Elem] =
SET_PARTIAL_CHOOSE [sort Elem]
then • \neg def choose(empty)
\forall S : Set \bullet def choose(S) \Rightarrow choose(S) is_in S
end
```

A definedness assertion, written 'def t', where t is a term, is a special kind of atomic formula: it holds if and only if the value of the term t is defined. For instance, in the above example, \neg def choose(empty) explicitly asserts that choose is undefined when applied to empty. Note that this axiom does not say anything about the definedness of choose applied to values other than empty, which means that choose may well be undefined on those values too. The second axiom of the above example asserts choose(S) is_in S under the condition def choose(S), to avoid undesired definedness induced by axioms, as explained in the previous section.

Note that if the two axioms of the above example were to be replaced by:

 $\forall S: Set \bullet \neg(S = empty) \Rightarrow choose(S) is_in S$

then we could conclude that choose(S) is defined when S is not equal to *empty*, but nothing about the undefinedness of choose(empty).

The domains of definition of partial functions can be specified exactly.

```
spec SET_PARTIAL_CHOOSE_2 [sort Elem] =
SET_PARTIAL_CHOOSE [sort Elem]
then \forall S : Set \bullet def choose(S) \Leftrightarrow \neg(S = empty)
\forall S : Set \bullet def choose(S) \Rightarrow choose(S) is_in S
end
```

In the above example, the domain of definition of the partial function choose is exactly specified by the axiom $def choose(S) \Leftrightarrow \neg(S = empty)$.

Loosely specified domains of definition may be useful.

```
spec NATURAL_WITH_BOUND_AND_ADDITION =
NATURAL_WITH_BOUND
then op __+?__: Nat \times Nat \rightarrow? Nat
\forall x, y : Nat
• def(x+?y) if x + y < max\_size
%{ x + y < max\_size implies both
x < max\_size and y < max\_size }%
• def(x+?y) \Rightarrow x+?y = x + y
```

end

In some cases, it is useful to loosely specify the domain of definition of a partial function, as illustrated in the above example for '+?', which is required to be defined for all arguments x and y such that $x + y < max_size$, but may well be defined on larger natural numbers as well. The point in loose specifications of definition domains is to avoid unnecessary constraints on the models of the specification. For instance, the above example does not exclude a model where '+?' is interpreted by a total function (which would then coincide with '+').²

Indeed, in some cases, specifying exactly domains of definition can be considered as overspecification. In most specifications, however, one would expect an exact specification of domains of definition, even for otherwise loosely specified functions (see, e.g., *choose* in SET_PARTIAL_CHOOSE_2).

Domains of definition can be specified more or less explicitly.

```
spec SET_PARTIAL_CHOOSE_3 [sort Elem] =
SET_PARTIAL_CHOOSE [sort Elem]
then • \neg def choose(empty)
\forall S : Set • \neg(S = empty) \Rightarrow choose(S) \text{ is_in } S
end
```

SET_PARTIAL_CHOOSE_3 specifies exactly the domain of definition of *choose*, but does this too implicitly, since some reasoning is needed to conclude that the above specification entails $def choose(S) \Leftrightarrow \neg(S = empty)$.

51

² In this example, it is essential to choose a new name '+?' for our partial addition operation. Otherwise, since '+' is (rightly) declared as a total operation in NATURAL_WITH_BOUND, the declaration **op** __+ -__: Nat × Nat \rightarrow ? Nat would be useless: the same name, same thing principle would lead to models with just one, total, addition operation.

To improve the clarity of specifications, it is in general advisable to specify definition domains as explicitly as possible, and SET_PARTIAL_CHOOSE_2 is somehow easier to understand than SET_PARTIAL_CHOOSE_3 (both specifications define the same class of models).

```
spec NATURAL_PARTIAL_PRE =
       NATURAL_ARITHMETIC
then op pre: Nat \rightarrow ? Nat
       • \neg def pre(0)
       \forall x : Nat \bullet pre(suc(x)) = x
```

end

In the above example, one can consider that the domain of definition of pre is (exactly) specified in an explicit enough way, since the first axiom states exactly that $pre(\theta)$ is undefined while the second one implies that pre is defined for all natural numbers of the form suc(x).

```
spec NATURAL_PARTIAL_SUBTRACTION_1 =
       NATURAL_PARTIAL_PRE
then op \_- \_: Nat \times Nat \rightarrow ? Nat
      \forall x, y : Nat
       • x - \theta = x
       • x - suc(y) = pre(x - y)
end
```

The above specification is perfect from a mathematical point of view, but is clearly not explicit enough, since there is no easy way to infer when x - y is defined. From a methodological point of view, the following alternative version is much better.

```
spec NATURAL_PARTIAL_SUBTRACTION =
       NATURAL_PARTIAL_PRE
then op \_-\_: Nat \times Nat \rightarrow ? Nat
       \forall x, y : Nat
        • def(x - y) \Leftrightarrow (y < x \lor y = x)
        • x - \theta = x
        • x - suc(y) = pre(x - y)
```

end

The above examples clearly demonstrate why the explicit specification of definition domains is generally advisable from a methodological point of view. However, they also indicate that this recommendation should not be applied in too strict a way, and that deciding whether a specification is explicit enough or not is to some extent a matter of taste.

Partial functions are minimally defined by default in free specifications.

```
spec LIST_SELECTORS_1 [sort Elem] =

LIST [sort Elem]

then free { ops head : List \rightarrow? Elem;

tail : List \rightarrow? List

\forall e : Elem; L : List

• head(cons(e, L)) = e

• tail(cons(e, L)) = L }
```

end

In the above example, the given axioms imply that *head* and *tail* are defined on lists of the form cons(e, L). The freeness constraint requires that these functions are minimally defined. Since the terms head(empty) and tail(empty) are not equated to any other term, the freeness constraint implies that these terms are undefined, and hence that the functions *head* and *tail* are undefined on *empty*. The situation here is similar to the fact that predicates hold minimally in models of free specifications (see Chap. 3, p. 38).

```
spec LIST_SELECTORS_2 [sort Elem] =

LIST [sort Elem]

then ops head : List \rightarrow? Elem;

tail : List \rightarrow? List

\forall e : Elem; L : List

• \neg def head(empty)

• \neg def tail(empty)

• head(cons(e, L)) = e

• tail(cons(e, L)) = L
```

end

The above specification LIST_SELECTORS_2 is an alternative to LIST_SELECTORS_1; both specifications define exactly the same class of models. However, LIST_SELECTORS_2 is clearly easier to understand and can be considered as technically simpler, since it involves no freeness constraint.

Operations like *head* and *tail* are usually called *selectors*, and CASL provides abbreviations to specify selectors in a very concise way, as we see next.

4.3 Partial Selectors and Constructors

Selectors can be specified concisely in datatype declarations, and are usually partial.

```
spec LIST_SELECTORS [sort Elem] =
free type List ::= empty | cons(head :? Elem; tail :? List)
```

The above free datatype declaration introduces, in addition to the constructors *empty* and *cons*, two partial selectors *head* and *tail* yielding the respective arguments of the constructor *cons*. Hence, this free datatype declaration with selectors has exactly the same effect as the ordinary free datatype declaration **free type** List ::= *empty* | *cons*(*Elem*; List), together with the operation declarations and axioms of LIST_SELECTORS_2 (i.e., LIST_ SELECTORS and LIST_SELECTORS_2 define exactly the same class of models). The following example is similar in spirit.

spec NATURAL_SUC_PRE = **free type** $Nat ::= 0 \mid suc(pre :? Nat)$

Selectors are usually total when there is only one constructor.

```
spec PAIR_1 [sorts Elem1, Elem2] =
free type Pair ::= pair(first : Elem1; second : Elem2)
```

While selectors are usually partial operations when there is more than one alternative in the corresponding datatype declaration, they can be total, and this is generally the case when there is only one constructor, as in the above example. The free datatype declaration entails in particular axioms asserting that *first* and *second* yield the respective arguments of the constructor *pair* (i.e., *first*(*pair*(e1, e2)) = e1 and *second*(*pair*(e1, e2)) = e2).

Constructors may be partial.

spec PART_CONTAINER [**sort** Elem] = **generated type** $P_{-Container} ::= empty | insert(Elem; P_{-Container})?$ **pred** $addable : Elem \times P_{-Container}$ **vars** $e, e' : Elem; C : P_{-Container}$ • $def insert(e, C) \Leftrightarrow addable(e, C)$ **pred** __is_in__: Elem × P_Container • $\neg(e \ is_in \ empty)$ • $(e \ is_in \ insert(e', C) \Leftrightarrow (e = e' \lor e \ is_in \ C)) \ if \ addable(e', C)$

end

The intention in the above example is to define a reusable specification of partial containers. The *insert* constructor is specified as a partial operation, defined if some condition on both the element e to be added and the container C to which the element is to be added holds. This condition is abstracted here in a predicate *addable*, so far left unspecified. Later on, instantiations of the PART_CONTAINER specification can be adapted to specific purposes by extending them with axioms defining *addable*.

The above generated datatype declaration abbreviates as usual the declaration of a sort $P_Container$, a constant constructor *empty*, and a partial constructor *insert* : $Elem \times P_Container \rightarrow ? P_Container$. It also entails the corresponding generatedness constraint.

4.4 Existential Equality

Existential equality requires the definedness of both terms as well as their equality.

```
spec NATURAL_PARTIAL_SUBTRACTION_2 =
NATURAL_PARTIAL_SUBTRACTION_1
then \forall x, y, z : Nat \bullet y - x \stackrel{e}{=} z - x \Rightarrow y = z
\% \{ y - x = z - x \Rightarrow y = z \text{ would be wrong,} def(y - x) \land def(z - x) \land y - x = z - x \Rightarrow y = z
is correct, but better abbreviated in the above axiom \}\%
```

end

An existential equation $t1 \stackrel{e}{=} t2$ is equivalent to $def(t1) \wedge def(t2) \wedge t1 = t2$, so it holds if and only if both terms t1 and t2 are defined and denote the same value. Existential equality $\stackrel{e}{=}$ ' is input as '=e='.

Note that while a trivial strong equation of the form t = t always holds, this is not the case for existential equations. For instance, the trivial existential equation $x - y \stackrel{e}{=} x - y$ does not hold, since the term x - y may be undefined.

In general consequences of undefinedness are undesirable. Hence a conditional equation of the form $t1 = t2 \Rightarrow t3 = t4$ is often wrong if t1 and t2 may be undefined, because the equality t3 = t4 would be implied when both t1 and t2 are undefined (since then the strong equation t1 = t2 would hold). The above specification provides a typical example of such a situation: $y - x = z - x \Rightarrow y = z$ would be wrong, since it would entail that any two

arbitrary values y and z are equal (it is enough to choose an x greater than y and z to make y - x and z - x both undefined).

Therefore, to avoid such undesirable consequences of undefinedness, it is advisable to use existential equations instead of strong equations in the premises of conditional equations involving partial operations. An alternative is to add the relevant definedness assertions explicitly to the equations in the premises.

Subsorting

Subsorts and supersorts are often useful in CASL specifications.

Many examples naturally involve subsorts and supersorts. CASL provides means for the declaration of a sort as a subsort of another one when the values of the subsort are regarded a special case of those in the other sort. The aim of this chapter is to discuss and illustrate how to handle subsorts and supersorts in CASL specifications.

5.1 Subsort Declarations and Definitions

Subsort declarations directly express relationships between carrier sets.

```
spec GENERIC_MONOID_1 [sort Elem] =

sorts Elem < Monoid

ops 1 : Monoid;

_*_: Monoid × Monoid → Monoid, assoc, unit 1
```

end

The above example declares the sort *Elem* to be a subsort of *Monoid*, or, symmetrically, the sort *Monoid* to be a supersort of *Elem*. Hence the specification GENERIC_MONOID_1 is quite similar to the specification GENERIC_MONOID given in Chap. 3, p. 30, the only difference being the use of a subsorting relation between *Elem* and *Monoid* instead of an explicit *inj* operation to embed values of sort *Elem* into values of sort *Monoid*.

In contrast to most other algebraic specification languages providing subsorting facilities, subsorts in CASL are interpreted by arbitrary *embeddings* between the corresponding carrier sets. In the above example, the subsort declaration *Elem < Monoid* induces an implicit (unnamed) embedding from the carrier of the sort *Elem* into the carrier of the sort *Monoid*. Thus the main difference between GENERIC_MONOID and GENERIC_MONOID_1 is that the embedding is explicit and named *inj* in GENERIC_MONOID while it is implicit in GENERIC_MONOID_1.

Note that interpreting subsorting relations by embeddings rather than inclusions does not exclude models where the (carrier of the) subsort happens to be a subset of (the carrier of) the supersort, and the embedding a proper inclusion. Embeddings are just slightly more general than inclusions, and technically not much more complex.

Operations declared on a sort are automatically inherited by its subsorts.

spec	Vehic	CLE =		
	NATU	RAL		
\mathbf{then}	\mathbf{sorts}	Car, Bicycle <	1	Vehicle
	\mathbf{ops}	max_speed	:	$Vehicle \rightarrow Nat;$
		weight	:	$Vehicle \rightarrow Nat;$
		engine_capacity	:	$Car \rightarrow Nat$

end

The above example introduces three sorts, *Car*, *Bicycle* and *Vehicle*, and declares both *Car* and *Bicycle* to be subsorts of *Vehicle*. A subsort declaration entails that any term of a subsort is also a term of the supersort, so here, any term of sort *Car* is also a term of sort *Vehicle*, and we can apply the operations *max_speed* and *weight* to it (and similarly for a term of sort *Bicycle*).

In other words, with the single declaration max_speed : $Vehicle \rightarrow Nat$, we get the effect of having declared also two other operations, max_speed : $Car \rightarrow Nat$ and max_speed : $Bicycle \rightarrow Nat$.¹

Obviously, operations that are only meaningful for some subsort should be defined at the appropriate level. This is the case here for the operation *engine_capacity*, which is only relevant for cars, and therefore defined with the appropriate profile exploiting the subsort *Car*.

¹ Strictly speaking, there is just one *max_speed* operation in the signature of VEHI-CLE. The difference between the kind of inheritance described here and operations actually declared on subsorts becomes important when writing symbol maps, see Chap. 7.

Inheritance applies also for subsorts that are declared afterwards.

spec MORE_VEHICLE = VEHICLE **then sorts** *Boat* < *Vehicle*

The order in which a subsort and an operation on the supersort are declared is irrelevant. In MORE_VEHICLE, we introduce a further subsort *Boat* of *Vehicle*, and as a consequence, we again get the effect of having both *max_speed* and *weight* available for boats, as was already the case for cars and bikes.

Subsort membership can be checked or asserted.

```
spec SPEED_REGULATION =

VEHICLE

then ops speed_limit : Vehicle \rightarrow Nat;

car_speed_limit, bike_speed_limit : Nat

\forall v : Vehicle

• v \in Car \Rightarrow speed_limit(v) = car_speed_limit

• v \in Bicycle \Rightarrow speed_limit(v) = bike_speed_limit
```

end

A subsort membership assertion, written ' $t \in s$ ', where t is a term and s is a sort, is a special kind of atomic formula: it holds if and only if the value of the term t is the embedding of some value of sort s. For instance, in the above example, $v \in Car$ holds if and only if v denotes a vehicle which is the embedding of a car value. Note that ' \in ' is input as 'in', but displayed as ' \in '.

Datatype declarations can involve subsort declarations.

The sequence of declarations:

sorts Car, Bicycle, Boat type Vehicle ::= sort Car | sort Bicycle | sort Boat

is equivalent to the declaration **sorts** Car, Bicycle, Boat < Vehicle. There may be some values of sort Vehicle which are not the embedding of any value of sort Car, Bicycle, or Boat. Intuitively, the above datatype declaration just means that Vehicle 'contains' the union (which may not be disjoint) of Car, Bicycle and Boat. Note that the subsorts used in the datatype declaration must already be declared beforehand.

The sequence of declarations:

```
sorts Car, Bicycle, Boat
generated type Vehicle ::= sort Car | sort Bicycle | sort Boat
```

is more restrictive, since the generatedness constraint implies that any value of the supersort *Vehicle* must be the embedding of some value of the declared subsorts *Car*, *Bicycle* and *Boat*. Intuitively, the above datatype declaration means that *Vehicle* 'is exactly' the union (which again may not be disjoint) of *Car*, *Bicycle* and *Boat*. In particular, this declaration prevents subsequent introduction of further subsorts (unless the values of the new subsorts are intended to correspond to some values of the already declared subsorts). For instance, if we were now to extend the above specification with **sorts** *Plane* < *Vehicle*, all values of sort *Plane* would have to correspond to *Car*, *Bicycle* or *Boat* values (which is presumably not what we were intending).

The sequence of declarations:

```
sorts Car, Bicycle, Boat
free type Vehicle ::= sort Car | sort Bicycle | sort Boat
```

entails the same generatedness constraint as in the previous example, and, moreover, the freeness constraint requires that there is no 'common' value in the subsorts of *Vehicle*. Intuitively, the above declaration means that *Vehicle* 'is exactly' the *disjoint* union of *Car*, *Bicycle* and *Boat*. This means in particular that the introduction of a further common subsort of both *Car* and *Boat* (say, **sorts** *Amphibious* < *Car*, *Boat*) is impossible.

Subsorts may also arise as classifications of previously specified values, and their values can be explicitly defined.

```
spec NATURAL_SUBSORTS =
NATURAL_ARITHMETIC
then pred even : Nat
• even(0)
• \neg even(1)
\forall n : Nat • even(suc(suc(n))) \Leftrightarrow even(n)
sort Even = {x : Nat • even(x)}
sort Prime = {x : Nat • 1 < x \land
\forall y, z : Nat • x = y * z \Rightarrow y = 1 \lor z = 1}
```

end

The subsort definition **sort** $Even = \{x : Nat \bullet even(x)\}$ is equivalent to the declaration of a subsort Even of Nat (i.e., **sorts** Even < Nat) together with the assertion $\forall x : Nat \bullet x \in Even \Leftrightarrow even(x)$.

The main advantage of defining the subsort *Even* in addition to the predicate *even* is that we may then use the subsort when declaring operations (e.g., **op** $times2 : Nat \rightarrow Even$) and variables.

The subsort definition for *Prime* above illustrates that it is not always necessary to introduce and define an explicit predicate characterizing the values of the subsort: the formula used in a subsort definition is not restricted to predicate applications. In fact whenever a (unary) predicate p on a sort s could be defined by **pred** $p(x : s) \Leftrightarrow f$ for some formula f, we may instead define **sort** $P = \{x : s \bullet f\}$, and use sort membership assertions $t \in P$ instead of predicate applications p(t), avoiding the introduction of the predicate p altogether.

The following example is a further illustration of subsort definitions. We declare a subsort *Pos* of *Nat* and ensure that values of *Pos* correspond to non-zero values of *Nat*. (Several alternative ways of specifying the sort *Pos* will be considered later in this section.)

spec POSITIVE = NATURAL_PARTIAL_PRE **then sort** $Pos = \{x : Nat \bullet \neg (x = 0)\}$

5.2 Subsorts and Overloading

It may be useful to redeclare previously defined operations, using the new subsorts introduced.

```
spec POSITIVE_ARITHMETIC =

POSITIVE

then ops 1 : Pos;

suc : Nat \rightarrow Pos;

-+-, -*-: Pos \times Pos \rightarrow Pos;

-+-: Pos \times Nat \rightarrow Pos;

-+-: Nat \times Pos \rightarrow Pos
```

end

Since, in POSITIVE, we have defined *Pos* as a subsort of *Nat*, all operations defined on natural numbers, like *suc*, '+' and '*' (see NATURAL_ARITHMETIC in Chap. 3, p. 38, which is extended into NATURAL_PARTIAL_PRE in Chap. 4, p. 52), are automatically inherited by *Pos* and can be applied to terms of sort *Pos*. However, according to their declarations, these operations, when applied to terms of sort *Pos*, yield results of sort *Nat*. To indicate that these results always correspond to values in the subsort *Pos*, it is necessary to explicitly *overload* these operations by similar ones with the appropriate profiles. This is

the aim of the first three lines of operation declarations in the above example. The last two operation declarations further overload '+' to specify that '+' also yields a result of sort Pos as soon as one of its arguments is a term of sort Pos.

It is quite important to understand that the above overloading declarations are enough to achieve the desired effect, and that *no axioms are necessary*. The fundamental rule is that, in models of CASL specifications with subsorting, embedding and overloading have to be compatible: embeddings commute with overloaded operations. This can be rephrased into the following intuitive statement: If terms look the same, then they have the same value in the same sort. Thus, in our example, the value of (1 + 1) is the same in Nat whatever the combination of the overloaded constant (1) and operation (+) is chosen, and there is no need for any axiom to ensure this, since this is implicit in the semantics of subsorting.

5.3 Subsorts and Partiality

A subsort may correspond to the definition domain of a partial function.

spec POSITIVE_PRE = POSITIVE_ARITHMETIC **then** op $pre : Pos \rightarrow Nat$

Since we have introduced the subsort Pos of non-zero natural numbers, it makes sense to overload the partial *pre* operation on *Nat* by a *total* one on *Pos*, as illustrated above, to emphasize the fact that indeed *pre* is a total operation on its definition domain. Note again that no further axiom is necessary, and that the semantics of subsorting will ensure that both the partial and total *pre* operations will give the same value when applied to the same non-zero value.²

Using subsorts may avoid the need for partial functions.

spec NATURAL_POSITIVE_ARITHMETIC = **free types** $Nat ::= 0 \mid sort Pos;$ Pos ::= suc(pre : Nat)

 $^{^2\,}$ This should not be confused with the same name, same meaning principle, which does not apply here: the total *pre* and the partial one have different profiles, and hence are just overloaded.

ops
$$1 : Pos = suc(0);$$

 $-+-: Nat \times Nat \rightarrow Nat, assoc, comm, unit 0;$
 $-*-: Nat \times Nat \rightarrow Nat, assoc, comm, unit 1;$
 $-+-, --*-: Pos \times Pos \rightarrow Pos;$
 $-+-: Pos \times Nat \rightarrow Pos;$
 $-+-: Nat \times Pos \rightarrow Pos$
 $\forall x, y : Nat$
• $x + suc(y) = suc(x + y)$
• $x * 0 = 0$
• $x * suc(y) = x + (x * y)$

end

It is indeed tempting to exploit subsorting to avoid the declaration of partial functions, as illustrated by the above NATURAL_POSITIVE_ARITHMETIC specification, which is an alternative to POSITIVE_PRE and avoids the introduction of the partial predecessor operation. Note that in the above example, we have fully used the facilities for defining free datatypes with subsorts, and in particular non-linear visibility for the declared sorts, which allows us to refer to the subsort *Pos* in the first line before defining it in the second one.

Avoiding the introduction of the partial predecessor operation has some drawbacks, since some previously well-formed terms (with defined values) have now become ill-formed, e.g., pre(pre(suc(1))), where pre(suc(1)) is a (well-formed) term of sort *Nat* but *pre* expects an argument of sort *Pos*. (The fact that pre(suc(1)) = 1 is a consequence of the specified axioms, and that 1 is of sort *Pos*, does not of course entail that pre(suc(1)) is of sort *Pos* too, since axioms are disregarded when checking for well-formedness.) See below for possible workarounds using explicit casts. Moreover, it is not always possible, or easy, to avoid the declaration of partial operations by using appropriate subsorts—just consider, e.g., subtraction on natural numbers.

As a last remark on this issue, the reader should be aware of the fact that, while overloading a partial operation on a supersort (say, *pre* on *Nat*) with a total one on a subsort (*pre* on *Pos*) is fine, overloading a total operation on a *supersort* with a partial one on a *subsort* forces the partial operation to be total, and hence, the latter would be better declared as total too.³

Casting a term from a supersort to a subsort is explicit and the value of the cast may be undefined.

³ Overloading a total *cons* on *List* with a partial *cons* on the subsort *OrderedList* would either lead to a total *cons* operation on *OrderedList*, or to an inconsistent specification, depending on how the definition domain of the partial *cons* is specified.

In CASL, a term of a subsort can always be considered as a term of any supersort, and embeddings are implicit. On the contrary, casting a term from a supersort to a subsort is explicit, and since casting is essentially a partial operation, the resulting casted term may not denote any value. Casting a term t to a sort s is written t as s. Consider the term pre(pre(suc(1)) as Pos) which is well-formed in the context of the NATURAL_POSITIVE_ARITHMETIC specification. This term does indeed denote a value, but the value is not a positive natural number, so the value of the term pre(pre(suc(1)) as Pos) as Pos is undefined.

Note that def(t as s) is equivalent to $t \in s$, for a well-formed term t of a supersort of s.

Supersorts may be useful when generalizing previously specified sorts.

```
spec INTEGER_ARITHMETIC_1 =
       NATURAL_POSITIVE_ARITHMETIC
then free type Int ::= sort Nat \mid -\_(Pos)
       ops \_+\_: Int \times Int \rightarrow Int, assoc, comm, unit <math>0;
             \_- \_: Int \times Int \rightarrow Int;
             \_*\_: Int \times Int \rightarrow Int, assoc, comm, unit 1;
             abs
                    : Int \rightarrow Nat
       \forall x : Int; n : Nat; p, q : Pos
        • suc(n) + (-1) = n
        • suc(n) + (-suc(q)) = n + (-q)
        • (-p) + (-q) = -(p+q)
        • x - \theta = x
        • x - p = x + (-p)
        • x - (-q) = x + q
        • \theta * (-q) = \theta
        • p * (-q) = -(p * q)
        • (-p) * (-q) = p * q
        • abs(n) = n
        • abs(-p) = p
```

end

The specification INTEGER_ARITHMETIC_1 extends NATURAL_POSITIVE_ ARITHMETIC and defines the sort *Int* as a supersort of the sort *Nat*. As a consequence, terms which have two parses in sort *Int*, depending whether the embedding from *Nat* to *Int* is applied to the arguments or to the result of overloaded operations, are required by the semantics of subsorting to have the same value for both parses, and they can therefore be used without explicit disambiguation. The situation would be quite different if one would be using a combination of NATURAL_ARITHMETIC and INTEGER_ARITHMETIC (see Chap. 3), say by extending both in a structured specification (see the next chapter for more details on structured specifications). In such a combination, a term such as $suc(\theta)$ would have two parses, one in sort *Nat* and one in sort *Int*; in the absence of any subsort declaration relating *Nat* and *Int*, and hence of any implicit embedding, this term would then be ambiguous, and would require explicit disambiguation to become a well-formed term.

Supersorts may also be used for extending the intended values by new values representing errors or exceptions.

end

The above specification SET_ERROR_CHOOSE is another variant of the various specifications of sets equipped with a partial choose function given in Chap. 4. This variant avoids the declaration of a partial function *choose* by using instead a supersort of *Elem*, namely *ElemError*, as the target sort of *choose*. The idea here is that values of *ElemError* which are not (embeddings of) values of *Elem* represent errors, e.g., the application of *choose* to the empty set. Note that to obtain the desired effect, it is necessary to explicitly state that *choose*(S) \in *Elem* when S is not the empty set; moreover, to make the term *choose*(S) *is_in* S well-formed, we have to explicitly overload the predicate *_is_in_*: *Elem*×*Set* provided by GENERATED_SET by a predicate *_is_in_*: *ElemError* × *Set* as shown above. This example demonstrates that avoiding partial functions by the use of 'error supersorts' is not as innocuous as it may seem, since in general one would need to enlarge the signatures considerably by adding all required overloadings.

```
spec SET_ERROR_CHOOSE_1 [sort Elem] =

GENERATED_SET [sort Elem]

then sorts Elem < ElemError

op choose : Set \rightarrow ElemError

\forall S : Set \bullet \neg(S = empty) \Rightarrow (choose(S) as Elem) is_in S
```

end

The specification Set_Error_Choose_1 above is a last attempt to avoid the use of partial functions; again, we introduce a supersort ElemError as

in SET_ERROR_CHOOSE, but to avoid the need for overloading the predicate is_in , we explicitly cast the term choose(S) in (choose(S) as Elem) $is_in S$. Note that when, for some value of S, (choose(S) as Elem) $is_in S$ holds, this implies that choose(S) as Elem is defined, and hence that $choose(S) \in Elem$ holds as well. This last version may seem preferable to the previous one. However, one should be aware that here, despite our attempt to avoid the use of partial functions, we rely on explicit casts, hence on terms that may not denote values: partiality has not been eliminated, the partial functions have merely been factorized as compositions of total functions with casting.
Structuring Specifications

Large and complex specifications are easily built out of simpler ones by means of (a small number of) specification-building operations.

In the previous chapters, we have focused attention on basic specifications and detailed how to use the various constructs of CASL to write meaningful, but relatively simple, specifications. The aim of this chapter is to discuss and illustrate how to assemble simple pieces of specifications into more complex, structured ones. In particular we explain how to extend specifications, make the union of several specifications, as well as how to rename or hide symbols when assembling specifications. Parametrization and instantiation of generic specifications are explained in the next chapter.

6.1 Union and Extension

Union and extension can be used to structure specifications.

```
spec LIST_SET [sort Elem] =

LIST_SELECTORS [sort Elem]

and GENERATED_SET [sort Elem]

then op elements_of__: List \rightarrow Set

\forall e : Elem; L : List

• elements_of empty = empty
```

• $elements_of \ cons(e, L) = \{e\} \cup elements_of \ L$

end

The above example shows how to make the union (expressed by 'and') of two specifications LIST_SELECTORS (see Chap. 4, p. 54) and GENERATED_SET (see Chap. 3, p. 35), and then further extend this union by an operation and some axioms (using 'then'). Union and extension are the most commonly used specification-building operations. In contrast with extension, whose purpose is to extend a given piece of specification by new symbols and axioms, union is generally used to combine two self-contained specifications. Union of specifications is obviously associative and commutative.

All symbols introduced by a specification are by default exported by it and visible in its extensions or in its unions with other specifications. (Variables are not considered as symbols, and never exported.) Remember also the 'same name, same thing' principle: in the above LIST_SET specification, it is therefore the same sort *Elem* which is used to construct both lists and sets.¹

 $Specifications may \ combine \ parts \ with \ loose, \ generated, \ and \ free \ interpretations.$

```
spec LIST_CHOOSE [sort Elem] =
       LIST_SELECTORS [sort Elem]
       SET_PARTIAL_CHOOSE_2 [sort Elem]
and
then ops elements_of_: List \rightarrow Set;
            choose : List \rightarrow? Elem
       \forall e : Elem: L : List
        • elements_of empty = empty
        • elements_of \ cons(e, L) = \{e\} \cup elements_of \ L
       • def choose(L) \Leftrightarrow \neg(L = empty)
        • choose(L) = choose(elements_of L)
end
spec SET_TO_LIST [sort Elem] =
       LIST_SET [ sort Elem ]
then op list_of_{-}: Set \to List
       \forall S : Set \bullet elements\_of(list\_of S) = S
```

\mathbf{end}

The specification LIST_CHOOSE is built as an extension of the union of LIST_SELECTORS and SET_PARTIAL_CHOOSE_2 (see Chap. 4, p. 50). This extension introduces an operation *elements_of* (as in LIST_SET) and a partial operation *choose*, which are defined by some axioms. In LIST_SELECTORS, lists are defined by a free datatype construct (with selectors), hence have a

¹ The constant *empty* is overloaded, since we have a constant *empty* : *List* for lists and a constant *empty* : *Set* for sets.

free interpretation. SET_PARTIAL_CHOOSE_2 is itself an extension of (SET_PARTIAL_CHOOSE which is an extension of) GENERATED_SET, where sets are defined by a generated datatype construct. However, note that as discussed in Chap. 3, p. 35, the apparently loose specification GENERATED_SET is in fact not so. Moreover, the *choose* partial function on sets is loosely defined in SET_PARTIAL_CHOOSE_2, and so is therefore also the *choose* partial function on lists defined in LIST_CHOOSE. It is easy to see that the operation *elements_of* is uniquely defined. The sort *Elem* has of course a loose interpretation.

Thus the specification LIST_CHOOSE combines parts with a free interpretation, parts with a generated interpretation, and parts with a loose interpretation. The situation is similar to that with LIST_SET (and SET_TO_LIST), where the operation *list_of* is loosely defined with the help of the operation *elements_of*.

6.2 Renaming

Renaming may be used to avoid unintended name clashes, or to adjust names of sorts and change notations for operations and predicates.

```
spec STACK [sort Elem] =
LIST_SELECTORS [sort Elem] with sort List \mapsto Stack,
ops cons \mapsto push\_onto\_,
head \mapsto top,
tail \mapsto pop
```

end

While the 'same name, same thing' principle has proven to be appropriate in numerous examples given in the previous chapters and above, it may still happen that, when combining specifications, this principle leads to unintended name clashes. An unintended name clash arises for instance when one combines two specifications that both export the same symbol (with the same profile in case of an operation or a predicate), while this symbol is not intended to denote the same 'thing' in the combination. In such cases, it is necessary to explicitly *rename* some of the symbols exported by the specifications put together in order to avoid the unintended name clashes.

When reusing a named specification, it may be convenient to rename some of its symbols; moreover, in the case of operation or predicate symbols, one may also change the style of notation. This is illustrated in the specification STACK above which is obtained as a renaming of the specification LIST_SELECTORS. (A renaming is introduced by the keyword 'with'.) First, the sort *List* is renamed into *Stack*, then the operation *cons* is renamed into a mixfix

operation *push_onto_*, and finally the selectors *head* and *tail* are renamed into *top* and *pop*, respectively. Note that ' \mapsto ' is input as '|->'.

The user only needs to indicate how symbols provided by the renamed specification are mapped to new symbols. A *signature morphism* is automatically deduced from this 'symbol map'. For instance, the signature morphism inferred from the symbol map specified in STACK maps the operation symbol *cons* : $Elem \times List \rightarrow List$ to the operation symbol *push_onto__* : $Elem \times Stack \rightarrow Stack$: not only the operation name is changed, but also its profile according to the renaming of *List* into *Stack*.

In a symbol map, one can qualify the symbol to be renamed by its kind, using the keywords **sort**, **op**, and **pred** (or their plural forms), as appropriate; this is illustrated in STACK above. Qualification in symbol maps is generally recommended since it improves their readability.

While it is possible to change the syntax of an operation or predicate symbol, as illustrated above for *cons* mapped to *push_onto__*, it is not possible to change the order of the arguments of the renamed operation or predicate.

In general, one does not need to rename all the symbols provided by the specification to be renamed. In the symbol map describing the intended renaming, it is indeed enough to mention only the symbols that change. By default, any symbol not explicitly mentioned is left unchanged (although its profile may be updated according to the renaming specified for some sorts). This is illustrated here in STACK where there is no need to rename the constant *empty*, which will therefore have the same name for both lists and stacks. However, the induced signature morphism maps the constant symbol *empty* : *List* into the constant symbol *empty* : *Stack*.

One can also explicitly rename a symbol to itself, say by writing ' $empty \mapsto empty$ ', or just mention it without providing a new name, as in 'with empty', which is equivalent to 'with $empty \mapsto empty$ '.

By default, overloaded symbols are renamed simultaneously. For instance, in INTEGER_ARITHMETIC_1 with $__+__ \mapsto plus$, all the five overloaded infix '+' operations exported by INTEGER_ARITHMETIC_1 (see Chap. 5, p. 64) are renamed into five *plus* operations, with a functional syntax and the appropriate profiles.

In general, it is possible to specifically rename one of some overloaded symbols, by specifying its profile in the symbol map. For instance, in LIST_SET with *empty* : *List* \mapsto *nil*, only the constant *empty* of sort *List* is renamed into *nil*, while the constant *empty* of sort *Set* remains unchanged. However, care is needed in the presence of subsorts, since the signature morphism induced by the specified symbol map should preserve the overloading relations associated with subsorts. For instance, if we attempt to only rename in the specification INTEGER_ARITHMETIC_1 the addition '+' of two positive numbers into *plus*, we merely obtain an ill-formed specification. Thus in the specification INTEGER_ARITHMETIC_1, all the five overloaded '+' operations must be renamed simultaneously, again into five overloaded symbols.

When combining specifications, origins of symbols can be indicated.

spec LIST_SET_1 [sort Elem] = LIST_SELECTORS [sort Elem] with empty, consand GENERATED_SET [sort Elem] with empty, {_}, _ \cup __ then op $elements_of_: List \rightarrow Set$ $\forall e : Elem; L : List$ • $elements_of \ empty = empty$

• $elements_of \ cons(e, L) = \{e\} \cup elements_of \ L$

end

Since, as explained above, 'with *empty*, *cons*' means 'with *empty* \mapsto *empty*, *cons* \mapsto *cons*', identity renaming can be used just to emphasize the fact that a given specification exports some symbols. This is illustrated in the specification LIST_SET_1 above, which is quite similar to LIST_SET, but for the fact that here we emphasize that LIST_SELECTORS exports in particular the operations *empty* and *cons*, and that GENERATED_SET exports in particular the operations *empty*, '{_}', and 'U'.

6.3 Hiding

Auxiliary symbols used in structured specifications can be hidden.

```
spec Natural_Partial_Subtraction_3 =
    Natural_Partial_Subtraction_1 hide suc, pre
end
```

spec Natural_Partial_Subtraction_4 =
 Natural_Partial_Subtraction_1
 reveal Nat, 0, 1, _+_, ___, ___, ____, ____<__</pre>

end

When writing large specifications, it is quite frequent to rely on auxiliary operations (and predicates) to specify the operations (and predicates) of interest. Once these are defined, the auxiliary operations are no longer needed, and are better removed from the exported signature of the specification, which should include only the symbols that were required to be specified. This is the purpose of the **hide** construct.

Consider for instance the specification NATURAL_PARTIAL_SUBTRACTION_ 1 given in Chap. 4, p. 52. Once addition and subtraction are defined, the two basic operations *suc* and *pre* are no longer needed (since suc(x) is more conveniently written x + 1, and similarly pre(x) is expressed by x - 1), and can therefore be hidden. This is illustrated by the specification NATURAL_ PARTIAL_SUBTRACTION_3 given above.

Depending on the relative proportion of symbols to be hidden or not, in some cases it may be more convenient to explicitly list the symbols to be exported by a specification rather than those to be hidden. The construct 'reveal' can be used for that purpose, and 'hide' and 'reveal' are just two symmetric constructs to achieve the same effect. The use of 'reveal' is illustrated in NATURAL_PARTIAL_SUBTRACTION_4 above, and the reader can convince himself that both NATURAL_PARTIAL_SUBTRACTION_3 and NATU-RAL_PARTIAL_SUBTRACTION_4 export exactly the same symbols. However, in this case the first specification is clearly more concise. A more convincing example of the use of 'reveal' is provided by the following example.

spec Partial_Order_2 = Partial_Order reveal pred $_ \leq _$

Similar rules to the ones explained for renaming apply to the **hide** and **reveal** constructs. One can qualify a symbol to be hidden or revealed by its kind (**sort**, **op** or **pred**), and by default, overloaded symbols are hidden (or revealed) simultaneously.

Note that hiding a sort entails hiding all the operations or predicates that use it in their profile. Similarly, revealing an operation or a predicate entails revealing all the sorts involved in its profile. For instance, in the specification PARTIAL_ORDER_2 above, revealing the predicate ' \leq ' entails revealing also the sort *Elem*.

As a consequence, hiding sorts should be used with care in the presence of subsorts. For instance, hiding the sort *Nat* in the specification POSITIVE given in Chap. 5, p. 61, leads to a specification of positive natural numbers with a sort *Pos* which has the expected carrier set, but without any operation or predicate available on it. Hiding the sort *Nat* in the specification POSITIVE_ARITHMETIC (see Chap. 5, p. 61) may seem more appropriate, but one should still note that the predicate '<' is no longer available in POSITIVE_ARITHMETIC hide sort *Nat*.

As a last remark, note that when convenient, **reveal** can be combined with a renaming of (some of) the exported symbols. For instance, in the above PARTIAL_ORDER_2 specification, we could have written '**reveal pred** $_ \leq _ \mapsto leq$ ' if, in addition to a restriction of the signature of PARTIAL_ ORDER, we wanted to rename the infix predicate '_ \leq __' into a predicate leqwith a functional notation.

6.4 Local Specifications

Auxiliary symbols can be made local when they do not need to be exported.

spec LIST_ORDER [TOTAL_ORDER with sort *Elem*, **pred** ___ < __] = LIST_SELECTORS [sort Elem] then local **op** $insert : Elem \times List \rightarrow List$ $\forall e, e' : Elem; L : List$ • insert(e, empty) = cons(e, empty)• insert(e, cons(e', L)) = cons(e', insert(e, L)) when e' < eelse cons(e, cons(e', L))within op $order : List \rightarrow List$ $\forall e : Elem; L : List$ • order(empty) = empty• order(cons(e, L)) = insert(e, order(L))end

In many cases, auxiliary symbols are introduced for immediate use, and they do not need to be exported by the specification where they are declared. Then the best is to limit the scope of the declarations of such auxiliary symbols by using the 'local ... within ...' construct. This is illustrated in the above specification LIST_ORDER, where the *insert* operation is introduced only for the purpose of the axiomatization of *order*. The declaration of *insert* has its scope limited to the part that follows 'within', and *insert* is therefore not exported by the specification LIST_ORDER.

It is generally advisable to ensure that auxiliary symbols are declared in local parts of specifications.

```
spec List_Order_Sorted
       [TOTAL_ORDER with sort Elem, pred _ < __] =
       LIST_SELECTORS [sort Elem]
then local
                pred __is_sorted : List
                \forall e, e' : Elem; L : List
                 • empty is_sorted
                 • cons(e, empty) is_sorted
                 • cons(e, cons(e', L)) is_sorted \Leftrightarrow
                      cons(e', L) is_sorted \land \neg(e' < e)
       within op order : List \to List
                \forall L: List \bullet order(L) is\_sorted
```

end

The specification LIST_ORDER_SORTED above is a variant of the specification LIST_ORDER illustrating again the use of the 'local ... within ...'

construct – this time to declare an auxiliary predicate. (Actually, the two specifications are not equivalent, since LIST_ORDER_SORTED is much looser and only requires that order(L) is a sorted list, but perhaps not with the same elements as L.)

Care is needed with local sort declarations.	
spec	WRONG_LIST_ORDER_SORTED
-	$[\text{TOTAL_ORDER with sort } Elem, \text{ pred } < _] =$
	LIST_SELECTORS [sort Elem]
then	local predis_sorted : List
	sort $SortedList = \{L : List \bullet L is_sorted\}$
	$\forall e, e' : Elem; \ L : List$
	• empty is_sorted
	• $cons(e, empty)$ is_sorted
	• $cons(e, cons(e', L))$ is_sorted \Leftrightarrow
	$cons(e', L)$ is sorted $\land \neg(e' < e)$
	within op $order: List \rightarrow SortedList$

end

Note that the above specification WRONG_LIST_ORDER_SORTED, which may at first glance be considered as a slight variant of LIST_ORDER_SORTED, is *ill-formed*: *order* is exported by WRONG_LIST_ORDER_SORTED, and hence all sorts occurring in its profile should also be exported, which cannot be, since the sort *SortedList* is auxiliary. So, if the specifier really intends to insist that the result sort of *order* is *SortedList*, this subsort should be exported, as shown below.

In fact the 'local ... within ...' construct abbreviates a combination of extension and explicit hiding. The specification LIST_ORDER_SORTED_2 above, for instance, abbreviates:

```
spec LIST_ORDER_SORTED_3
  [TOTAL_ORDER with sort Elem, pred _ < __] =
  LIST_SELECTORS [sort Elem]
then {
    pred __is_sorted : List
        \forall e, e' : Elem; L : List
        empty is\_sorted
        cons(e, empty) is\_sorted
        cons(e, cons(e', L)) is\_sorted \Leftrightarrow
        cons(e', L) is\_sorted \land \neg(e' < e)
then sort SortedList = {L : List • L is\_sorted}
        op order : List → SortedList
    }
    hide __is\_sorted
```

end

The main advantage of using the 'local ... within ...' construct is that hiding the symbols introduced in the local part is left implicit. The convenience of this generally outweighs the danger of overlooking a locally-declared sort that is needed for the profile of an exported symbol. In any case, CASL allows both styles, and users can simply choose the one they prefer.

6.5 Named Specifications

Naming a specification allows its reuse.

It is in general advisable to define as many named specifications as felt appropriate, since this improves the reusability of specifications: a named specification can easily be reused by referring to its name.

Not only do the names serve as abbreviations when writing specifications, they also make it easy for readers to notice reuse. Moreover, when the name of a specification is aptly chosen, e.g., NATURAL_ARITHMETIC, readers may well be able to guess its signature – and perhaps even the specified axioms – from the name itself. (In Chap. 9, we shall see how named specifications and other items can be collected in libraries, and particular versions of them made available for use over the Internet.)

References to named specifications are particularly convenient for specifications structured using unions and extensions, where verbatim insertion of unnamed specifications would tend to obscure the structure. When needed, the signature of a referenced specification can be adjusted through appropriate combinations of renaming and hiding (although this should not often be necessary, provided that auxiliary symbols are made local, as explained in the previous section).

Generic Specifications

Making a specification generic (when appropriate) improves its reusability.

As mentioned in the previous chapter, naming specifications is a good idea. In many cases, however, datatypes are naturally *generic*, having sorts, operations, and/or predicates that are deliberately left loosely specified, to be determined when the datatype is used. For instance, datatypes of lists and sets are generic regarding the sort of elements. *Generic specifications* allow the genericity of a datatype to be made explicit by declaring *parameters* when the specification is named: in the case of lists and sets, there is a single parameter that simply declares the sort *Elem*.¹ A fitting argument specification has to be provided for each parameter of a generic specification.

The aim of this chapter is to discuss and illustrate how to define generic specifications and instantiate them. We have seen plenty of simple examples of generic specifications and instantiations in the previous chapters. In more complicated cases, however, explicit fitting symbol maps may be required to determine the exact relationship between parameters and arguments in instantiations, and so-called *imports* should be separated from the bodies of generic specifications.

¹ Generic specifications are also useful to ensure loose coupling between several named specifications, replacing an explicit extension by a parameter including only the necessary symbols and their required properties. This is illustrated in the Steam-Boiler Control System case study, see Chap. 13.

[©] IFIP International Federation for Information Processing 2004

7.1 Parameters and Instantiation

Parameters are arbitrary specifications.

Any specification, named or not, can be used as the parameter of a generic specification. Commonly, the parameter is a rather trivial specification consisting merely of a single sort declaration, as in most of the examples given in the previous chapters, e.g.:

```
spec GENERIC_MONOID [sort Elem] = %{ See Chap. 3, p. 30 }%
```

```
spec LIST_SELECTORS [sort Elem] = %{ See Chap. 4, p. 54 }%
```

However, the parameter can also be a more complex, possibly structured, specification, as in:

```
spec LIST_ORDER [TOTAL_ORDER with sort Elem, pred ___ < __] = %{ See Chap. 6, p. 73 }%
```

Recall that 'with' requires the signature of the specification to include the listed symbols; here, in fact, the signature of TOTAL_ORDER does not contain any further symbols, so those are all the symbols that have to be supplied when instantiating LIST_ORDER.

The argument specification of an instantiation must provide symbols corresponding to those required by the parameter.

```
spec LIST_ORDER_NAT = LIST_ORDER [NATURAL_ORDER]
```

The correspondence between the symbols provided by the argument specification and those required by the parameter can be described by a fitting symbol map or left implicit when it is not ambiguous, which is often the case.

In the above example, the argument specification NATURAL_ORDER provides the sort *Nat*, the operation symbols θ and *suc*, and the binary predicate symbol '<'. Hence this specification indeed provides symbols corresponding to those required by the parameter specification TOTAL_ORDER and the correspondence can be left implicit because the argument NATURAL_ORDER has only single symbols of the right kind. (The coincidence of the predicate symbol in the parameter and argument is irrelevant here.)

How to describe explicit fitting symbol maps and when they can be omitted is detailed later in this section.

The argument specification of an instantiation must ensure that the properties required by the parameter hold.

spec NAT_WORD = GENERIC_MONOID [NATURAL]

A (fitting) signature morphism from the signature of the parameter part to the signature of the argument specification is automatically deduced, taking into account the explicitly specified fitting symbol map if any (the situation here is quite similar to a renaming, where a signature morphism is deduced from a symbol map). The instantiation is *defined* if all models of the argument specification, when reduced along the induced fitting signature morphism, provide models of the parameter part. In particular the symbols provided by the argument specification must have the properties, if any, specified in the parameter for their counterparts. When this is the case, we get not only a signature morphism, but also a (fitting) specification morphism from the argument specification to the parameter specification.²

In the above NAT_WORD example, since the parameter of GENERIC_ MONOID is trivial, it is obvious that the instantiation is defined.

The effect of the instantiation is to make the union of the argument specification and of the (non generic equivalent of the) generic specification, renamed according to the induced fitting signature morphism. In particular, a side-effect of the instantiation is to rename the symbols of the generic specification according to the fitting signature morphism induced by the instantiation. In our NAT_WORD example, the operation symbol $inj : Elem \rightarrow Monoid$ is renamed into $inj : Nat \rightarrow Monoid$, while the operation symbols '1' and '*' are left unchanged (as well as the sort Monoid). Thus, the specification NAT_ WORD abbreviates the following specification:

```
NATURAL and { NON_GENERIC_MONOID with Elem \mapsto Nat }.
```

When convenient, the instantiation can be completed by a renaming, as illustrated in the following variant of NAT_WORD.

```
spec NAT_WORD_1 =
GENERIC_MONOID [NATURAL]
with Monoid \mapsto Nat_Word
```

end

In the case of the specification LIST_ORDER_NAT above, checking the definedness of the instantiation corresponds to a non-trivial proof obligation. The instantiation is defined since the predicate '<' provided by NATURAL_ ORDER is indeed a total ordering relation, hence the properties required by

 $^{^2\,}$ Note that consistency is entirely orthogonal to definedness: a defined instantiation may be consistent or not.

TOTAL_ORDER are fulfilled, even if there is no syntactic correspondence between the axioms given in TOTAL_ORDER and those in NATURAL_ORDER.

There must be no shared symbols between the argument specification and the body of the instantiated generic specification.

spec This_Is_Wrong = Generic_Monoid [Monoid]

The intention in the above example may have been to specify monoids of monoids. However, the above instantiation is ill-formed since the sort *Monoid* and the operation symbols '1' and '*' are shared between the body of the generic specification GENERIC_MONOID and the argument specification MONOID.

Section 7.3 provides useful hints on how to structure generic specifications in order to avoid as far as possible undesirable clashes of symbols in instantiations. A correct specification of monoids of monoids is provided in Sect. 7.2, p. 86.

In instantiations, the fitting of parameter symbols to identical argument symbols can be left implicit.

```
spec GENERIC_COMMUTATIVE_MONOID [sort Elem] = GENERIC_MONOID [sort Elem]
```

then ...

When the parameter and the argument have symbols in common, these parameter symbols are implicitly taken to fit directly to the corresponding argument symbols. Thus it is *never* necessary to make explicit that a symbol is mapped identically. In the above example, for instance, the parameter specification of GENERIC_MONOID is exactly the same as the argument specification in its instantiation, so the fitting can be left implicit.

The fitting of parameter sorts to unique argument sorts can also be left implicit.

When the argument specification has only a single sort, the fitting of all parameter sorts to that sort is obvious, and can again be left implicit, as illustrated earlier by the NAT_WORD specification. Of course, this does not apply the other way round: if the parameter has a single sort (which is often

the case in practice) but the argument specification has more than one sort, the parameter sort could be mapped to any of the argument sorts, so the fitting symbol map has to be made explicit – except when the parameter sort is identical to one of the argument sorts, as previously explained, or when the fitting of sorts can be implied from the fitting of other symbols, as explained below.

Fitting of operation and predicate symbols can sometimes be left implicit too, and can imply fitting of sorts.

spec LIST_ORDER_POSITIVE = LIST_ORDER [POSITIVE]

Fitting of operation and predicate symbols can imply fitting of sorts. For instance, when a parameter predicate symbol is fitted to an argument predicate symbol whose profile involves different sorts, this implies that the parameter sorts involved have to be fitted to the corresponding sorts in the argument specification.

This is illustrated in the above LIST_ORDER_POSITIVE specification. In a first step, the fitting of the parameter sort *Elem* to one of the argument sorts *Nat* and *Pos* provided by the specification POSITIVE (see Chap. 5, p. 61) may seem ambiguous. However, no explicit fitting of symbols is necessary here, since the argument specification provides only one binary predicate symbol, and the fitting of the corresponding binary predicate symbol of the parameter specification to it entails the fitting of the sort *Elem* to the sort *Nat* (Again, the coincidence of the predicate symbol in the parameter and argument is irrelevant here.)

As may be clear by now, the exact rules for when the fitting between parameter and argument symbols can be left implicit are quite sophisticated. It seems best to make the intended fitting explicit whenever it is not completely obvious, using the notation for fitting arguments illustrated in the following examples.

The intended fitting of the parameter symbols to the argument symbols may have to be specified explicitly.

spec NAT_WORD_2 = GENERIC_MONOID [NATURAL_SUBSORTS fit $Elem \mapsto Nat$]

The correspondence between the symbols required by the parameter and those provided by the argument specification can be made explicit using socalled *fitting symbol maps*. For instance, the above NAT_WORD_2 specification, which differs from NAT_WORD only regarding the presence of subsorts of *Nat*, is obtained as an instantiation of GENERIC_MONOID, fitting the parameter part '**sort** *Elem*' to the NATURAL_SUBSORTS specification. The mapping between the parameter sort *Elem* and the sort *Nat* provided by NATURAL_SUBSORTS is described by the fitting symbol map 'fit $Elem \mapsto Nat'$.

A generic specification may have more than one parameter.

spec PAIR [sort Elem1] [sort Elem2] =
free type Pair ::= pair(first : Elem1; second : Elem2)

Using several parameters is merely a notational convenience, since they are equivalent to their union. For instance, the above PAIR specification is nothing but a variant of the specification PAIR_1 with just one parameter 'sorts *Elem1*, *Elem2*' defined in Chap. 4, p. 54.

Note that writing:

```
spec HOMOGENEOUS_PAIR_1 [sort Elem] [sort Elem] =
free type Pair ::= pair(first : Elem; second : Elem)
```

merely defines pairs of values of the same sort, and HOMOGENEOUS_PAIR_1 is (equivalent to and) better defined as follows:

```
spec HOMOGENEOUS_PAIR [sort Elem] =
free type Pair ::= pair(first : Elem; second : Elem)
```

since the two parameters in HOMOGENEOUS_PAIR_1 are equivalent to just one 'sort *Elem*' parameter.

From a methodological point of view, it is generally advisable to use as many parameters as convenient: the part of the specification that is intended to be specialized at instantiation time is better split into logically coherent units, each one corresponding to a parameter. Consider for instance:

```
spec TABLE [sort Key] [sort Val] = ...
```

Here, using two parameters in TABLE emphasizes that Key and Val are logically distinct entities which can be specialized as desired independently of each other.

Instantiation of generic specifications with several parameters is similar to the case of just one parameter.

```
spec PAIR_NATURAL_COLOR =
PAIR [NATURAL_ARITHMETIC] [COLOR fit Elem_2 \mapsto RGB]
```

In the above example, the first parameter 'sort Elem1' of PAIR is instantiated by NATURAL_ARITHMETIC, which exports only one sort Nat, hence no explicit fitting symbol map is necessary. The second parameter 'sort Elem2' of PAIR is instantiated by COLOR: in this case a fitting symbol map must be provided, since COLOR exports two sorts, RGB and CMYK.

Using the specification PAIR_1 would require us to write:

spec PAIR_NATURAL_COLOR_1 = PAIR_1 [NATURAL_ARITHMETIC and COLOR fit $Elem1 \mapsto Nat, Elem2 \mapsto RGB$]

which clearly demonstrates the benefit of using two parameters as in PAIR instead of just one as in PAIR_1.

When parameters are trivial ones (i.e., just one sort), one can always avoid explicit fitting maps. Consider for instance the following alternative to PAIR_NATURAL_COLOR:

spec PAIR_NATURAL_COLOR_2 =
 PAIR [sort Nat] [sort RGB]
and NATURAL_ARITHMETIC and COLOR

This may be convenient when the argument specification exports several sorts. Compare for instance:

```
spec PAIR_POS =
    HOMOGENEOUS_PAIR [sort Pos] and INTEGER_ARITHMETIC_1
```

with:

```
spec PAIR_POS_1 =
HOMOGENEOUS_PAIR [INTEGER_ARITHMETIC_1 fit Elem \mapsto Pos]
```

Note that the instantiation:

HOMOGENEOUS_PAIR_1 [NATURAL] [COLOR fit $Elem \mapsto RGB$]

is ill-formed, since it entails mapping the sort *Elem* to both *Nat* and *RGB*.

More generally, care is needed when the several parameters of a generic specification share some symbols, which in general is not advisable.

As a last remark, note that it is easy to specialize a generic specification with several parameters, using a 'partial instantiation', as in the following version of TABLE:

```
spec My_TABLE [sort Val] =
TABLE [NATURAL_ARITHMETIC] [sort Val]
```

where we still have a parameter for the values to be stored, but have decided that the keys are natural numbers.

Composition of generic specifications is expressed using instantiation.

spec SET_OF_LIST [**sort** Elem] = GENERATED_SET [LIST_SELECTORS [**sort** Elem] **fit** $Elem \mapsto List$]

The above generic specification SET_OF_LIST describes sets of lists of arbitrary elements, and is obtained by an instantiation of the generic specification GENERATED_SET, whose parameter 'sort *Elem*' is instantiated by the specification LIST_SELECTORS, itself trivially instantiated. Since the (trivially instantiated) specification LIST_SELECTORS exports two sorts *Elem* and *List*, it is of course necessary to specify, in the instantiation of GENERATED_SET, the fitting symbol map from the parameter sort *Elem* to the argument sort *List*.

Note especially that the following specification:

```
spec MISTAKE [sort Elem] =
   GENERATED_SET [LIST_SELECTORS [sort Elem]]
```

does *not* provide sets of lists of elements: The sort *Elem* in the parameter of GENERATED_SET is mapped by the identity fitting symbol map to the sort *Elem* provided by the instantiation of the generic specification LIST_ SELECTORS [sort *Elem*], rather than to the sort *List.*³ Thus MISTAKE just provides sets of arbitrary elements *and* lists of arbitrary elements. If this was indeed the desired effect, then one should rather write instead:

spec SET_AND_LIST [sort Elem] =
 GENERATED_SET [sort Elem] and LIST_SELECTORS [sort Elem]

As illustrated by SET_OF_LIST, composition of generic specifications is fairly easy in CASL. Note however that this composition is achieved by means of appropriate instantiations (some possibly trivial), and that CASL does not provide higher-order genericity.

It may be worth mentioning that the following composition of generic specifications is ill-formed:

```
spec THIS_IS_STILL_WRONG =
GENERIC_MONOID [ GENERIC_MONOID [ sort Elem ]
fit Elem \mapsto Monoid ]
```

³ However, the situation would be different if the parameter of GENERATED_SET had been, e.g., 'sort Val', since then the absence of an explicit fitting symbol map would have led to an ambiguity: in that case the specifier would have to specify whether the sort Val is to be mapped to *Elem* or to *List*.

The above instantiation is ill-formed since the sort *Monoid* and the operation symbols '1' and '*' are shared between the body of the generic specification GENERIC_MONOID and the argument specification GENERIC_MONOID [sort *Elem*] (where this time the generic specification GENERIC_MONOID is trivially instantiated). The next section provides (p. 86) a correct specification of monoids of monoids.

7.2 Compound Symbols

Compound sorts introduced by a generic specification get automatically renamed on instantiation, which avoids name clashes.

```
spec LIST_REV [sort Elem] =
      free type List[Elem] ::= empty
                                 cons(head :? Elem; tail :? List[Elem])
      ops \_++\_: List[Elem] \times List[Elem] \rightarrow List[Elem],
                           assoc, unit empty;
            reverse : List[Elem] \rightarrow List[Elem]
      \forall e : Elem; L, L1, L2 : List[Elem]
       • cons(e, L1) + L2 = cons(e, L1 + L2)
       • reverse(empty) = empty
       • reverse(cons(e, L)) = reverse(L) + cons(e, empty)
end
```

```
spec LIST_REV_NAT = LIST_REV [NATURAL]
```

A compound sort is a sort of the form '*Name*[*Name1*,...,*NameN*]'. In the specification LIST_REV, we introduce a compound sort *List*[*Elem*] to denote lists (of arbitrary elements), instead of the simple sort *List* used in the previous examples. When the specification LIST_REV is instantiated as in LIST_REV_ NAT, the translation induced by the (implicit) fitting symbol map is applied to the component *Elem* also where it occurs in *List*[*Elem*], providing a sort *List*[*Nat*]. Thus, compound sorts can be seen as a convenient way of implicitly completing the instantiation by an appropriate renaming of the (compound) sorts introduced by the generic specification.

```
spec
     TWO\_LISTS =
      LIST_REV [NATURAL] %% Provides the sort List[Nat]
      LIST_REV [COLOR fit Elem \mapsto RGB] %% Provides the sort List[RGB]
and
```

Using a compound sort *List*[*Elem*] proves particularly useful in the above example Two_LISTS, where we make the union of two distinct instantiations of LIST_REV. If we had used an ordinary sort *List*, then an unintentional name clash would have arisen,⁴ and we would have to complete each instantiation by an explicit renaming of the sort *List*.

Note that in the specification Two_LISTS, we have two sorts *List*[*Nat*] and List[RGB], hence two overloaded constants *empty* (one of each sort), which may need disambiguation when used in terms. (How to disambiguate terms is explained in Chap. 3, p. 31.)

Similarly, we have overloaded operation symbols cons, head, tail, ++, and *reverse*, but in general their context of use in terms will be enough to disambiguate which one is meant.

```
TWO\_LISTS\_1 =
spec
      LIST_REV [INTEGER_ARITHMETIC_1 fit Elem \mapsto Nat]
      LIST_REV [INTEGER_ARITHMETIC_1 fit Elem \mapsto Int]
and
```

Since the specification INTEGER_ARITHMETIC_1 provides three sorts Nat, *Pos*, and *Int*, an explicit fitting symbol map is needed in the above instantiations, which provide the sorts List[Nat] and List[Int]. Note that the subsorting relation Nat < Int does not entail List[Nat] < List[Int], but of course this can be added if desired in an extension by a subsorting declaration.

Using compound sorts, we can now easily specify monoids of monoids.

```
spec MONOID_C [sort Elem] =
       sort Monoid[Elem]
                    : Elem \rightarrow Monoid[Elem];
       ops inj
              1
                    : Monoid [Elem];
              -*: Monoid [Elem] \times Monoid [Elem] \rightarrow Monoid [Elem],
                      assoc, unit 1
       \forall x, y : Elem \bullet inj(x) = inj(y) \Rightarrow x = y
```

end

```
spec MONOID_OF_MONOID [sort Elem] =
      MONOID_C [MONOID_C [sort Elem] fit Elem \mapsto Monoid[Elem]]
```

The instantiation in MONOID_OF_MONOID is now correct, since the use of a compound sort *Monoid* [*Elem*] ensures there is no clash of symbols between the body of the instantiated generic specification and the argument specification.

⁴ And the specification Two_LISTS would have been inconsistent, due to the same name, same thing principle and the fact that *List* is defined by a free type construct.

Compound symbols can also be used for operations and predicates.

```
spec LIST_REV_ORDER [TOTAL_ORDER] =
       LIST_REV [sort Elem]
               op insert : Elem \times List[Elem] \rightarrow List[Elem]
then local
               \forall e, e' : Elem; L : List[Elem]
                • insert(e, empty) = cons(e, empty)
                • insert(e, cons(e', L)) = cons(e', insert(e, L)) when e' < e
                                           else cons(e, cons(e', L))
       within op order[\_ < \_] : List[Elem] \rightarrow List[Elem]
               \forall e : Elem; L : List[Elem]
                • order[\_ < \_](empty) = empty
                • order[\_ < \_](cons(e, L)) = insert(e, order[\_ < \_](L))
end
spec LIST_REV_WITH_TWO_ORDERS =
       LIST_REV_ORDER
        [INTEGER_ARITHMETIC_ORDER fit Elem \mapsto Int, ... < \_ \mapsto ... < \_]
        %% Provides the sort List[Int] and the operation order[-<-]
and
      LIST_REV_ORDER
        [INTEGER_ARITHMETIC_ORDER fit Elem \mapsto Int, \_ < \_ \mapsto \_ > \_]
        %% Provides the sort List[Int] and the operation order[->-]
then % implies
```

 $\forall L : List[Int] \bullet order[_- < _-](L) = reverse(order[_- > _-](L))$ end

The above example illustrates the use of compound identifiers for operation symbols, and the same rules apply to predicate symbols. While in most cases using compound identifiers for sorts will be sufficient, in some situations it is also convenient to use them for operation or predicate symbols, as done here for order[... < ...]. When LIST_REV_ORDER is instantiated, not only does the sort List[Elem] get renamed (here, to List[Int]), but also the operation symbol order[... < ...], according to the fitting symbol map corresponding to the instantiation. If we had not used a compound identifier for the order operation, then an unintentional name clash would have arisen. Note that on the other hand we rely on the same name, same thing principle to ensure that the sorts List[Int] provided by each of the two instantiations are the same, which indeed is what we want for this example.

Of course we do not bother to use a compound identifier for the *insert* operation symbol. This operation being local, it is not exported by LIST_REV_ORDER and cannot be the source of unintentional name clashes in instantiations.

7.3 Generic Specifications with Imports

Parameters should be distinguished from references to fixed specifications that are not intended to be instantiated.

spec LIST_WEIGHTED_ELEM [**sort** *Elem* **op** *weight* : *Elem* \rightarrow *Nat*] **given** NATURAL_ARITHMETIC =

LIST_REV [sort *Elem*]

```
then op weight : List[Elem] \rightarrow Nat
```

```
\forall e : Elem; L : List[Elem]
```

- weight(empty) = 0
- weight(cons(e, L)) = weight(e) + weight(L)

end

In the above example, we specialize lists of arbitrary elements to lists of elements equipped with a *weight* operation, which is then overloaded by a *weight* operation on lists. Therefore we specify that the generic specification LIST_WEIGHTED_ELEM has for parameter a specification extending the 'given' specification NATURAL_ARITHMETIC by a sort *Elem* and an operation symbol *weight*. Thereby the intention is to emphasize the fact that only the sort *Elem* and the operation *weight* are intended to be specialized when the specification LIST_WEIGHTED_ELEM is instantiated, and not the 'fixed part' NATURAL_ARITHMETIC. In CASL, the specifications listed after the 'given' keyword are called *imports*. One could have written instead:

spec List_Weighted_Elem

[NATURAL_ARITHMETIC then sort $Elem \text{ op } weight : Elem \rightarrow Nat$] = ...

but the latter, which is correct, misses the essential distinction between the part which is intended to be specialized and the part which is 'fixed' (since, by definition, the parameter is the part which has to be specialized).

Note also that omitting the '**given** NATURAL_ARITHMETIC' clause would make the declaration:

```
spec LIST_WEIGHTED_ELEM [sort Elem op weight : Elem \rightarrow Nat ] = ...
```

ill-formed, since the sort *Nat* is not available.

To summarize, the 'given' construct is useful to distinguish the 'true' parameter from the part which is 'fixed'. Both the parameter and the body of the generic specification extend what is provided by the imports (i.e., by the specifications listed after the 'given' keyword), whose exported symbols are therefore available.

Argument specifications are always implicitly regarded as extension of the imports.

spec LIST_WEIGHTED_PAIR_NATURAL_COLOR = LIST_WEIGHTED_ELEM [PAIR_NATURAL_COLOR **fit** $Elem \mapsto Pair$, $weight \mapsto first$]

The instantiation specified in LIST_WEIGHTED_PAIR_NATURAL_COLOR is correct since the fitting symbol map is the identity on all the symbols exported by the 'fixed part' NATURAL_ARITHMETIC (which happens here to be included in the argument specification PAIR_NATURAL_COLOR). More generally, the argument specification is always regarded as an extension of the imports, and the fitting symbol map should be the identity on all symbols provided by these imports. This is illustrated in the next example:

```
spec LIST_WEIGHTED_INSTANTIATED =
LIST_WEIGHTED_ELEM [sort Value op weight : Value \rightarrow Nat]
```

Here we rely on a rather trivial instantiation (whose purpose is merely to illustrate our point) where the fitting symbol map can be omitted since no ambiguity arises and where the argument specification 'sort Value op weight : Value \rightarrow Nat' is well-formed because it is regarded as an extension of the imports of LIST_WEIGHTED_ELEM (i.e., as an extension of NATURAL_ ARITHMETIC), which implies that the sort Nat is available.

Imports are also useful to prevent ill-formed instantiations.

```
spec LIST_LENGTH [sort Elem] given NATURAL_ARITHMETIC =

LIST_REV [sort Elem]

then op length : List[Elem] \rightarrow Nat

\forall e : Elem; L : List[Elem]

• length(empty) = 0

• length(cons(e, L)) = length(L) + 1

then %implies
```

```
\forall L : List[Elem] \bullet length(reverse(L)) = length(L)
```

end

The specification LIST_LENGTH needs the sort *Nat* and the usual arithmetic operations provided by NATURAL_ARITHMETIC to specify the *length* operation. In this case it is clear that the imports have nothing to do with the (trivial) parameter of LIST_LENGTH. The reason to specify NATURAL_ARITHMETIC as an import is that this will make instantiations of LIST_LENGTH similar to the following one well-formed.

90 7 Generic Specifications

```
spec List_Length_Natural =
    List_Length [Natural_Arithmetic]
```

To understand this point, consider the following variant of LIST_LENGTH:

The specification WRONG_LIST_LENGTH is fine as long as one does not need to instantiate it with NATURAL_ARITHMETIC as argument specification. The instantiation WRONG_LIST_LENGTH [NATURAL_ARITHMETIC] is ill-formed since some symbols of the argument specification are shared with some symbols of the body (and not already occurring in the parameter) of the instantiated generic specification, which is wrong, as already explained p. 80. Of course the same problem will occur with any argument specification which provides, e.g., the sort *Nat*.

In generic specifications, auxiliary required specifications should be imported rather than extended.

As illustrated by the above examples, one should remember the following essential point. Since an instantiation is ill-formed as soon as there are some shared symbols between the argument specification and the body of the generic specification, when designing a generic specification, it is generally advisable to turn auxiliary required specifications (such as NATURAL-ARITHMETIC for LIST_LENGTH) into imports, and generic specifications of the form 'F[X] = SP then ...' are better written 'F[X] given SP = ...' to allow the instantiation 'F[SP]'.

7.4 Views

Views are named fitting maps, and can be defined along with specifications.

```
view INTEGER_AS_TOTAL_ORDER :

TOTAL_ORDER to INTEGER_ARITHMETIC_ORDER =

Elem \mapsto Int, \_ < \_ \mapsto \_ < \_
```

```
view INTEGER_AS_REVERSE_TOTAL_ORDER :
TOTAL_ORDER to INTEGER_ARITHMETIC_ORDER =
Elem \mapsto Int, \ \_ < \_ \mapsto \_ > \_
```

```
spec LIST_REV_WITH_TWO_ORDERS_1 =
    LIST_REV_ORDER [view INTEGER_AS_TOTAL_ORDER]
and LIST_REV_ORDER [view INTEGER_AS_REVERSE_TOTAL_ORDER]
then %implies
    ∀L : List[Int] • order[._<_](L) = reverse(order[._>_](L))
```

end

A view is nothing but a convenient way to name a specification morphism (induced by a symbol map) from a (parameter) specification to an (argument) specification. The rules regarding the omission of 'evident' symbol maps in explicit fittings apply to views too. A view proves particularly useful when the same instantiation (with the same fitting symbol map) is intended to be used several times: naming a specification morphism once and for all makes its reuse easier. Once a view is defined, as e.g. INTEGER_AS_TOTAL_ORDER [view INTEGER_AS_TOTAL_ORDER], where the keyword 'view' makes it clear that the argument is not merely a named specification with an implicit fitting map, which would be written differently.

Since a view is defined only when the given symbol map induces a specification morphism (i.e., all models of the target specification, when reduced along the signature morphism induced by the given symbol map, provide models of the source specification), it may be convenient to use views just to explicitly document the existence of some specification morphisms, even when these are not intended to be used in any instantiation. For instance, the view INTEGER_AS_TOTAL_ORDER can be seen as the assertion that INTEGER_ ARITHMETIC_ORDER indeed specifies '<' to be a total ordering relation, and would therefore make sense even without being used later on in instantiations.

Views can also be generic.

view LIST_AS_MONOID [**sort** Elem]: MONOID **to** LIST_REV [**sort** Elem] = $Monoid \mapsto List[Elem], 1 \mapsto empty, ...* ... \mapsto ... + +...$

A view can be generic, being then defined with some parameters (as illustrated above in the LIST_AS_MONOID view) and possibly some imports. The reader should be aware that, in a generic view, the target specification (here, the trivially instantiated specification LIST_REV) is not interpreted as such, but as the body of a generic specification with the same parameters and imports as the view. (The source specification is on the contrary interpreted exactly as provided.)

The above example illustrates again the use of a view as a 'proof obligation', asserting that lists (equipped with the '++' operation) form a monoid.

Specifying the Architecture of Implementations

Architectural specifications impose structure on implementations, whereas specification-building operations only structure the text of specifications.

As explained in the previous chapters, the specification of a complex system may be fairly large and should be structured into coherent, easy to grasp, pieces. CASL provides a number of specification-building operations to achieve this, as detailed in Chap. 6. Moreover, generic specifications, described in Chap. 7, provide pieces of specification that are easy to reuse in different contexts, where they can be adapted as desired by instantiating them.

Specification-building operations and generic specifications are useful to structure the text of the specification of the system under consideration. However, the models of a structured specification have no more structure than do those of a flat, unstructured, specification. Indeed, most examples given in the previous chapters could have been structured differently, with the same meaning (i.e., with the same models). Structured specifications are usually adequate at the requirements stage, where the focus is on the expected overall properties of the system under consideration.

In contrast, the aim of architectural specifications is to prescribe the intended architecture of the implementation of the system. Architectural specifications provide the means for specifying the various components from which the system will be built, and describing how these components are to be assembled to provide an implementation of the system of interest. At the same time, they allow the task of implementing a system to be split into independent, clearly-specified sub-tasks. Thus, architectural specifications are essential at the design stage, where the focus is on how to factor the implementation of the system into components.

The aim of this chapter is to discuss and illustrate both the role of architectural specifications and how to express them in CASL.

The idea underlying architectural specifications is that eventually in the process of systematic development of modular software from specifications, components are implemented as software modules in some chosen programming language. However, this step is beyond the scope of specification formalisms, so in CASL and in this chapter we identify components with models (and with functions from models to models, in the case of generic components). The modular structure of the software under development, as described by an architectural specification, is therefore captured here simply as an explicit, structural way to build CASL models.

The examples in this chapter are artificially simple.

Architectural specifications, and more generally component-oriented approaches, are intended for relatively large systems. In this chapter, however, we have to rely on simple small examples to illustrate and explain CASL architectural specification concepts and constructs. After reading this chapter, the reader is encouraged to study Chap. 13, which provides realistic examples of the use of architectural specifications. A more detailed account of the rationale behind architectural specifications in the context of formal software development by stepwise refinement can be found in [11].

The following structured specifications will be referred to later in this chapter when illustrating CASL architectural specifications:

```
spec COLOR = \% As defined in Chap. 3, p. 37 }%
spec NATURAL_ORDER = \% { As defined in Chap. 3, p. 38 }%
spec NATURAL_ARITHMETIC = \% { As defined in Chap. 3, p. 38 }%
spec ELEM = sort Elem
spec CONT [ELEM] =
       generated type Cont[Elem] ::= empty | insert(Elem; Cont[Elem])
       preds __is_empty : Cont[Elem];
                \_is\_in\_: Elem \times Cont[Elem]
                choose : Cont[Elem] \rightarrow ? Elem;
       ops
                delete: Elem \times Cont[Elem] \rightarrow Cont[Elem]
       \forall e, e' : Elem; C : Cont[Elem]
        • empty is_empty
        • \neg insert(e, C) is_empty
        • \neg e \ is in \ empty
        • e \text{ is\_in insert}(e', C) \Leftrightarrow (e = e' \lor e \text{ is\_in } C)
        • def choose(C) \Leftrightarrow \neg C is_empty
        • def choose(C) \Rightarrow choose(C) is_in C
        • e \text{ is_in } delete(e', C) \Leftrightarrow (e \text{ is_in } C \land \neg(e = e'))
```

end

```
spec CONT_DIFF [ELEM] =
        CONT [ELEM]
then op diff : Cont[Elem] \times Cont[Elem] \rightarrow Cont[Elem]
        \forall e : Elem; C, C' : Cont[Elem]
         • e \text{ is_in } diff(C, C') \Leftrightarrow (e \text{ is_in } C \land \neg(e \text{ is_in } C'))
end
spec Req = Cont_Diff [Natural_Order]
spec FLAT_REQ =
        free type Nat ::= 0 \mid suc(Nat)
        pred \_ < \_ : Nat \times Nat
        generated type Cont[Nat] ::= empty \mid insert(Nat; Cont[Nat])
        preds __is_empty : Cont[Nat];
                  \_is\_in\_: Nat \times Cont[Nat]
                  choose : Cont[Nat] \rightarrow ? Nat;
        ops
                  delete: Nat \times Cont[Nat] \rightarrow Cont[Nat];
                  diff: Cont[Nat] \times Cont[Nat] \rightarrow Cont[Nat]
        \forall e, e' : Nat; C, C' : Cont[Nat]
         • \theta < suc(e)
         • \neg(e < \theta)
         • suc(e) < suc(e') \Leftrightarrow e < e'
         • empty is_empty
         • \neg insert(e, C) is_empty
         • \neg e \ is\_in \ empty
         • e \text{ is\_in insert}(e', C) \Leftrightarrow (e = e' \lor e \text{ is\_in } C)
         • def choose(C) \Leftrightarrow \neg C is_empty
         • def choose(C) \Rightarrow choose(C) is_in C
         • e \text{ is_in } delete(e', C) \Leftrightarrow (e \text{ is_in } C \land \neg(e = e'))
         • e \text{ is_in } diff(C, C') \Leftrightarrow (e \text{ is_in } C \land \neg(e \text{ is_in } C'))
end
```

8.1 Architectural Specifications

Let's assume in the following that REQ describes our requirements about the system to be implemented. First, note that both REQ and FLAT_REQ have the same models, which illustrates our point about the fact that the CASL specification-building operations are merely facilities to structure the text of specifications into coherent units.

An architectural specification consists of a list of unit declarations, specifying the required components, and a result part, indicating how they are to be combined.

arch spec SYSTEM = units N: NATURAL_ORDER; C: CONT [NATURAL_ORDER] given N;

D: Cont_Diff [Natural_Order] given C

result D

The SYSTEM architectural specification is intended to prescribe a specific architecture for implementing the system specified by REQ.

The first part, introduced by the keyword **units**, indicates that we require the implementation of our system to be made of three components N, C, and D. The second part, introduced by the keyword **result**, indicates that the component D provides the desired implementation.

Each component is provided with its specification. The line:

N : Natural_Order

declares a component N specified by NATURAL_ORDER, which means simply that N should be a model of NATURAL_ORDER.

The line:

C: CONT [NATURAL_ORDER] given N

declares a component C which, given the previously declared component N, provides a model of CONT [NATURAL_ORDER]. It is essential to understand that the component C must *expand* the assumed component N into a model of CONT [NATURAL_ORDER], which means that C reduced to the signature of NATURAL_ORDER must be equal to N. This property reflects the fact that a software module is supposed to use what it is given exactly as supplied, without altering it.

Similarly, the line:

D: CONT_DIFF [NATURAL_ORDER] given C

declares a component D which, given the component C, expands it into a model of CONT_DIFF [NATURAL_ORDER].

The final result is therefore simply D. (More complex examples of result expressions will be illustrated in examples below.)

As in the rest of CASL, visibility is linear in architectural specifications, meaning that any component must be declared before being used (e.g., the component N should be declared before being referred to by 'given N' in the declaration of the component C in the architectural specification SYSTEM).

Component names (such as N, C, and D in SYSTEM) are *local* to the architectural specification where they are declared, and are not visible outside it.

There can be several distinct architectural choices for the same requirements specification.

```
arch spec System_1 =
units N : NATURAL_ORDER;
    CD : CONT_DIFF [NATURAL_ORDER] given N
result CD
```

The architectural specifications SYSTEM and SYSTEM_1 both provide models of REQ. However, the former insists on an implementation made of three components, while the latter insists on an implementation made of two components. Thus the architectural specification SYSTEM_1 corresponds to a different architectural choice for implementing our REQ specification. Of course, further design for implementing the component CD of SYSTEM_1 may lead to splitting this implementation task exactly as in SYSTEM above. However, there are also other possibilities, including for instance an architectural design where we would split the task of implementing CD into two different tasks, one for implementing containers with all their operations (including *diff*) except *delete*, the other for implementing *delete* by means of *diff* and other operations.

Each unit declaration listed in an architectural specification corresponds to a separate implementation task.

For instance, in the architectural specification SYSTEM, the task of providing a component D expanding C and implementing CONT_DIFF [NATURAL_ ORDER] is independent from the tasks of providing implementations N of NATURAL_ORDER and C of CONT [NATURAL_ORDER] given N. Hence, when providing the component D, one cannot make any further assumption on how the component C is (or will be) implemented, besides what is expressly ensured by its specification.

To understand this, let us consider again the requirements specification REQ (or its variant FLAT_REQ). Among its models, there is one where containers are implemented by sorted lists (in increasing order, without repetitions), and in this model we can choose to implement *diff* by the following algorithm:

 $\begin{aligned} diff(L,L') &= nil \ when \ L = nil \\ else \ L \ when \ L' = nil \\ else \ insert(head(L), diff(tail(L), L')) \ when \ head(L) < head(L') \\ else \ diff(tail(L), tail(L')) \ when \ head(L) = head(L') \\ else \ diff(L, tail(L')) \end{aligned}$

In this model, however, we rely on knowledge about the implementation of containers to decide how to implement diff – which is fine, since both are simultaneously implemented in the same component. In contrast, in the architectural specification SYSTEM, we request that containers are to be implemented in the component C while diff is to be provided by a separate component D. Imposing that the component D can be developed independently of the component C means that for D it is no longer possible to implement diffas sketched above, since this specific implementation choice may not be compatible with an independently chosen realization for C (where containers may be implemented by bags, for instance). Hence an implementation of *diff* in the component D can only rely on the operations provided by C (e.g., *choose* and delete); this may turn out to be less efficient for some particular realization of C, but should be compatible with any independently chosen realization for C(bags, for instance). In the case of the architectural specification SYSTEM_1, since both containers and the *diff* operation are implemented in the same component CD, we can of course decide to implement containers by ordered lists without repetitions and *diff* as sketched above.

Thus the component D should expand any given implementation C of CONT [NATURAL_ORDER] and provide an implementation of CONT_DIFF [NATURAL_ORDER], which is tantamount to providing a generic implementation G of CONT_DIFF [NATURAL_ORDER] which takes the particular implementation of CONT_[NATURAL_ORDER] as a parameter to be expanded. Then we obtain D by simply applying G to C.

Genericity here arises from the independence of the developments of C and D, rather than from the desire to build multiple implementations of CONT_DIFF [NATURAL_ORDER] using different implementations of CONT [NATURAL_ORDER]. This is reflected by the fact that G is left implicit in the architectural specification SYSTEM.

A unit can be implemented only if its specification is a conservative extension of the specifications of its given units.

For instance, the component D can exist only if the specification CONT_ DIFF [NATURAL_ORDER] is a conservative extension of CONT [NATURAL_ ORDER], i.e., if any model of the latter specification can be expanded into a model of the former one, which is indeed the case here. Similarly, the component C can exist since CONT [NATURAL_ORDER] is a conservative extension of NATURAL_ORDER. Consider now the following variant of CONT_DIFF [NATURAL_ORDER] and the associated variant of the architectural specification SYSTEM.

spec CONT_DIFF_1 = CONT [NATURAL_ORDER] then op diff : $Cont[Nat] \times Cont[Nat] \rightarrow Cont[Nat]$ $\forall x, y : Nat; C, C' : Cont[Nat]$ • diff(C, empty) = C• diff(empty, C') = empty• diff(insert(x, C), insert(y, C')) = insert(x, diff(C, insert(y, C'))) when x < y else diff((C, C') when x = y else diff(insert(x, C), C')• $x \text{ is_in } diff(C, C') \Leftrightarrow (x \text{ is_in } C \land \neg(x \text{ is_in } C'))$ end arch spec INCONSISTENT —

arch spec INCONSISTENT =
units N : NATURAL_ORDER;
C : CONT [NATURAL_ORDER] given N;
D : CONT_DIFF_1 given C
result D

The specification CONT_DIFF_1 is consistent (has some models, for instance sorted lists, in increasing order, without repetitions), but is not a conservative extension of CONT [NATURAL_ORDER] (since, for instance, a model of CONT [NATURAL_ORDER] where containers are realized by arbitrary lists, possibly with repetitions, cannot be expanded into a model of CONT_DIFF_1 – in that case, the last two axioms are contradictory). As a consequence, in the architectural specification INCONSISTENT, the specification of the component D is inconsistent, since no component can expand *all* implementations C of CONT [NATURAL_ORDER] into models of CONT_DIFF_1. The architectural specification INCONSISTENT is therefore itself inconsistent.

To summarize, architectural specifications not only prescribe the intended architecture of the implementation of the system, but they also ensure that the specified components can be developed independently of each other (which imposes a certain degree of genericity for these components).

8.2 Generic Components

Genericity of components can be made explicit in architectural specifications.

```
arch spec System_G =
```

units N: NATURAL_ORDER;

```
F : NATURAL_ORDER \rightarrow CONT [NATURAL_ORDER];
```

 $G: \text{ CONT [NATURAL_ORDER]} \rightarrow \text{CONT_DIFF [NATURAL_ORDER]}$ result G[F[N]]

The architectural specification SYSTEM_G is a variant of SYSTEM; here we choose to specify the second and third components as explicit *generic* components.

The line:

F: Natural_Order \rightarrow Cont [Natural_Order]

declares a generic component F. Given any component implementing (i.e., model of) NATURAL_ORDER, F should expand it into an implementation of CONT [NATURAL_ORDER]. The models of the generic-component specification NATURAL_ORDER \rightarrow CONT [NATURAL_ORDER] are functions that map any model of NATURAL_ORDER to a model of CONT [NATURAL_ORDER]. These functions are required to be persistent, meaning that the result model expands the argument model.

The third component G is also specified as a generic component: given any implementation of CONT [NATURAL_ORDER], G should expand it into an implementation of CONT_DIFF [NATURAL_ORDER].

Hence the whole system is obtained by the composition of applications G[F[N]], as described in the result part. In CASL, such combinations of components are called *unit terms*. (More complex examples of unit terms will be illustrated in examples below.)

The component C of SYSTEM corresponds to the application F[N] in SYSTEM_G, and similarly the component D in SYSTEM corresponds to G[C], i.e., to G[F[N]] in SYSTEM_G.

The models of a specification of the form $SP1 \rightarrow SP2$ are generic components GC that should always expand their argument into a model of the target specification. This only makes sense as long as the signature of the target specification contains the signature of SP1. This is why in CASL, SP2 is always considered as an implicit extension of SP1, and $SP1 \rightarrow SP2$ abbreviates $SP1 \rightarrow \{ SP1 \text{ then } SP2 \}$.¹ Moreover, since the generic component GC

¹ When *SP2* is already defined as an extension of *SP1*, as it is the case for instance here for CONT_DIFF [NATURAL_ORDER], *SP2* is equivalent to *SP1* then *SP2*.

should expand any model of SP1, the specification $SP1 \rightarrow SP2$ is consistent (i.e., has some models) if and only if the specification SP1 then SP2 is a conservative extension of SP1. Forgetting this fact is a potential source of inconsistent specifications of generic components in architectural specifications. For instance, the specification CONT [NATURAL_ORDER] \rightarrow CONT_DIFF_1 is inconsistent, for the reasons explained at the end of the previous section.

A generic component may be applied to an argument richer than required by its specification.

arch spec System_A =

units NA : NATURAL_ARITHMETIC;

F : NATURAL_ORDER \rightarrow CONT [NATURAL_ORDER];

 $G : \text{CONT} [\text{NATURAL_ORDER}] \rightarrow \text{CONT_DIFF} [\text{NATURAL_ORDER}]$ result G[F[NA]]

The above architectural specification SYSTEM_A is a variant of SYSTEM_G. Here we require a component NA implementing the specification NATURAL_ARITHMETIC, instead of a component N implementing NATURAL_ORDER as in SYSTEM_G (perhaps because we know that such a component is already available in some collection of previously-implemented components.)

The generic component F requires a component fulfilling the specification NATURAL_ORDER, but can of course be applied to a richer argument, as in F[NA]. A similar reasoning applies to G.

More generally, a generic component can be applied to any component (or to any unit term) that can be reduced along some morphism to an argument of the required 'type' (i.e., to a model of the required specification). When necessary, a fitting symbol map can be used to describe the correspondence between the symbols provided by the argument and those expected by the generic component. We do not detail here the technicalities related to these fitting symbol maps, since they are quite similar to those used in instantiations of generic specifications and the notations are the same.

As a last remark, note that, similarly to what happens when instantiating a generic specification by an argument specification, when a generic component is applied to an argument richer than required, the extra symbols are kept in the result. Hence the result of the architectural specification SYSTEM_A above contains also the interpretations of the arithmetic and ordering operations on natural numbers, as they are provided by the component NA. This means in particular that the implementations described by SYSTEM_A have a larger signature than the ones described by SYSTEM_G.

Specifications of components can be named for further reuse.

```
unit spec CONT_COMP = ELEM \rightarrow CONT [ELEM]

unit spec DIFF_COMP = CONT [ELEM] \rightarrow CONT_DIFF [ELEM]

arch spec SYSTEM_G1 =

units N : NATURAL_ORDER;

F : CONT_COMP;

G : DIFF_COMP

result G[F[N]]
```

In the above example, we give the name CONT_COMP to the specification (of generic components) ELEM \rightarrow CONT [ELEM]. Similarly, we give the name DIFF_COMP to the specification CONT [ELEM] \rightarrow CONT_DIFF [ELEM]. Then both named specifications can be reused in the architectural specification SYSTEM_G1 which is similar to the architectural specification SYSTEM_G.

In the architectural specification SYSTEM_G1, we use again the fact that the generic component F can be applied to richer arguments than models of ELEM (and similarly for G). Since ELEM is more general (has more models) than NATURAL_ORDER, there are potentially fewer possibilities for implementing the generic component specified by CONT_COMP (which should be compatible with any model of ELEM) than there are for implementing the generic component specified by NATURAL_ORDER \rightarrow CONT [NATURAL_ ORDER] (which only needs to be compatible with models of NATURAL_ ORDER; a similar argument holds for DIFF_COMP). As a consequence, the architectural specifications SYSTEM_G and SYSTEM_G1 do not describe the same implementations of the requirements specification REQ.

Both named and unnamed specifications can be used to specify components.

unit spec DIFF_COMP_1 = CONT [ELEM] \rightarrow { op diff : Cont[Elem] \times Cont[Elem] \rightarrow Cont[Elem] $\forall e : Elem; C, C' : Cont[Elem]$ • $e \text{ is_in diff}(C, C') \Leftrightarrow$ ($e \text{ is_in } C \land \neg(e \text{ is_in } C')$) }

So far we have always used named (structured) specifications to specify components. unnamed specifications can be used as well, as illustrated by the above variant DIFF_COMP_1 of DIFF_COMP. Here, for the sake of the example, we directly specify the *diff* operation instead of referring to the named specification CONT_DIFF. Remember that in a specification of a generic component of the form $SP1 \rightarrow SP2$, SP2 is always considered as an implicit extension of SP1, which explains why the above example is well-formed.

Specifications of generic components should not be confused with generic specifications.

Generic specifications naturally give rise to specifications of generic components, which can be named for later reuse, as illustrated above by CONT_ COMP. However, the reader should not confuse a generic specification (which is nothing other than a piece of specification that can easily be adapted by instantiation) with the corresponding specification of a generic component: the latter cannot be instantiated, it is the specified *generic component* which gets *applied* to suitable components.

Conservative extensions of the form 'spec SP2 = SP1 then SP' also naturally give rise to specifications of generic components of the form $SP1 \rightarrow SP2$, as illustrated by DIFF_COMP above.

A generic component may be applied more than once in the same architectural specification.

arch spec OTHER_SYSTEM = units N: NATURAL_ORDER; C: COLOR; F: CONT_COMP result F[N] and F[C fit $Elem \mapsto RGB]$

The above architectural specification requires a component N specified by NATURAL_ORDER, a component C specified by COLOR, and a generic component F specified by CONT_COMP. Then, as described by the result part, the desired system is obtained by applying F to N and applying F to C(in this case, an explicit fitting symbol map is necessary, since COLOR exports two sorts RGB and CMYK). Finally both application results are combined, which is expressed by 'and'.

Apart from '**free**', all specification-building operations for structured specifications have natural counterparts at the level of components, which are expressed using the same keywords.² The reader should remember that specification-building operations work with specifications defining classes of

² The situation is however a bit different with specification extensions, which lead to specifications of generic components, as explained above, or to specifications of components expanding a given component, as illustrated in the previous section.
models (e.g., union of specifications, denoted by 'and'), while in architectural specifications we work with individual models (corresponding to components, as is the case here in OTHER_SYSTEM where 'and' is used to combine the two components F[N] and F[C fit $Elem \mapsto RGB]$).

Hence renaming and hiding also have natural counterparts at the level of components. For instance, remember that the implementations described by SYSTEM_A have a larger signature than the implementations described by SYSTEM_G. It is however easy to modify the result part of SYSTEM_A if what we really want are implementations with the same signature as the implementations described by SYSTEM_G: one has just to hide the extra symbols resulting from the component NA as follows:

result G[F[NA]] hide 1, $_+_$, $_-*_$

or:

```
result G[F[NA \text{ hide } 1, \_+\_, \_*\_]]
```

Symbol maps used in renaming and hiding at the level of components follow the same rules as symbol maps used in renaming and hiding at the level of structured specifications (see Chap. 6).

Several applications of the same generic component is different from applications of several generic components with similar specifications.

```
arch spec OTHER_SYSTEM_1 =

units N : NATURAL_ORDER;

C : COLOR;

FN : NATURAL_ORDER \rightarrow CONT [NATURAL_ORDER];

FC : COLOR \rightarrow CONT [COLOR fit Elem \mapsto RGB]

result FN[N] and FC[C]
```

The above architectural specification OTHER_SYSTEM_1 is a variant of OTHER_SYSTEM. However, in OTHER_SYSTEM, we insist on choosing one implementation for containers in the generic component F, and then we apply it twice, first to a component N implementing NATURAL_ORDER, and then to a component C implementing COLOR. In contrast, in OTHER_SYSTEM_1, we may choose two different implementations for containers, one for containers of natural numbers in the component FN and another one for containers of colors in the component FC.

The architectural specifications OTHER_SYSTEM and OTHER_SYSTEM_1 are therefore similar but clearly different. Neither is better than the other: each corresponds to a different architectural decision, and selecting one rather than the other is a matter of architectural design. Components that are more widely reusable tend to have less efficient implementations, in general. (Here

the fact that RGB has only three values might be exploited in FC to give a more space-efficient representation of containers than is possible for FN.)

Generic components may have more than one argument.

```
unit spec SET_COMP = ELEM \rightarrow GENERATED_SET [ELEM]

spec CONT2SET [ELEM] =

CONT [ELEM] and GENERATED_SET [ELEM]

then op elements_of_-: Cont[Elem] \rightarrow Set

\forall e : Elem; C : Cont[Elem]

• elements_of empty = empty

• elements_of insert(e, C) = {e} \cup elements_of C

end

arch spec Arch_CONT2SET_NAT =

units N : NATURAL_ORDER;

C : CONT_COMP;

S : SET_COMP;

F : CONT [ELEM] \times GENERATED_SET [ELEM]

\rightarrow CONT2SET [ELEM]
```

```
result F[C[N]][S[N]]
```

The architectural specification ARCH_CONT2SET_NAT requires a component N implementing NATURAL_ORDER, a generic component C implementing CONT_COMP, i.e., containers, and a generic component S implementing SET_COMP, i.e., sets. Then it further requires a generic component F that, given any pair of compatible models X of CONT [ELEM] and Y of GENERATED_SET [ELEM], expands them into a model of CONT2SET [ELEM].

Models X and Y are said to be *compatible* if they share a common interpretation for all symbols they have in common. Here the only symbol they have in common is the sort *Elem*, so the compatibility condition means that X and Y have the same carrier set for *Elem*. Compatibility is a natural condition, since it is obviously necessary that X and Y have a common interpretation of their common symbols, otherwise they cannot be both expanded to the same more complex component.

The result is then produced by applying F to the pair obtained by applying C to N and S to N. Here the pair of arguments C[N] and S[N] are obviously compatible, since their common symbols (the sort *Nat* equipped with the operations θ and *suc*) all come from the *same* component N which provides their interpretation, which is expanded (hence cannot be modified) in C[N] and in S[N], thus compatibility is guaranteed.

Open systems can be described by architectural specifications using generic unit expressions in the result part.

```
arch spec Arch_Cont2Set =

units C: Cont_COMP;

S: Set_COMP;

F: Cont [ELEM] × GENERATED_SET [ELEM]

\rightarrow CONt2Set [ELEM]

result \lambda X: ELEM • F [C[X]] [S[X]]

arch spec Arch_Cont2Set_Used =

units N : NATURAL_ORDER;

CSF: arch spec Arch_Cont2Set

result CSF [N]
```

So far our example architectural specifications have described 'closed', stand-alone systems where all components necessary to build the desired system were declared in the architectural specification of interest. In CASL, it is however possible to describe 'open' systems, i.e., systems made of some components that would require further components to provide a 'closed' system. This is illustrated by the architectural specification ARCH_CONT2SET which describes a system with a generic component C implementing containers, a generic component S implementing sets, and a generic component F that expands them to provide an implementation of the operation *elements_of*. The result part is therefore a generic structured component, i.e., an 'open' system, which, given any component X implementing ELEM, provides a system built by applying F to the pair made of the applications of C to X and of S to X. In CASL, ' λ ' is input as 'lambda'.

As illustrated by ARCH_CONT2SET_USED, we can then describe a 'closed' system made of a component N implementing NATURAL_ORDER, and of an 'open' system CSF specified by ARCH_CONT2SET, which is then applied to N in the result part.

8.3 Writing Meaningful Architectural Specifications

In the previous sections we have already pointed out potential sources of inconsistent specifications of components. Another issue which deserves some attention when designing an architectural specification is *compatibility* between components (or, more generally, unit terms) that are to be combined together, either by 'and', or by fitting them to a generic component with multiple arguments. When components are to be combined, it is best to check that any shared symbol originates from the same non-generic component.

```
arch spec Arch_Cont2Set_Nat_1 =

units N: NATURAL_ORDER;

C: CONT_COMP;

S: SET_COMP;

G: { CONT [ELEM] and GENERATED_SET [ELEM] }

\rightarrow CONT2SET [ELEM]

result G[C[N] and S[N] fit Cont[Elem] \mapsto Cont[Nat]]
```

The architectural specification ARCH_CONT2SET_NAT_1 is a variant of ARCH_CONT2SET_NAT where, instead of declaring a generic component F with two arguments, we now declare a generic component G with a single argument, which must be a model of the specification { CONT [ELEM] and GENERATED_SET [ELEM] }, obtained as the union of the two (trivially instantiated) specifications of containers and sets.

As a consequence, to obtain the desired system, in the result part we apply the generic component G to the combination (denoted by 'and') of C applied to N and of S applied to N.³ This combination makes sense only if both C[N]and S[N] share the same interpretation of their common symbols. Here their common symbols (the sort *Nat* equipped with the operations θ and *suc*) all come from the *same* component N which provides their interpretation, which is expanded (hence cannot be modified) in C[N] and in S[N], thus compatibility is guaranteed.

There is a clear analogy here between the application of the generic component F with multiple arguments in ARCH_CONT2SET_NAT and the combination of C[N] and S[N] in ARCH_CONT2SET_NAT_1: in both cases the result is meaningful because we can trace shared symbols like the sort *Nat* and the operations θ and *suc* to a single component N introducing them.

Let us emphasize again that compatibility is a natural requirement: since each unit declaration corresponds to a separate implementation task (and hence each unit subterm to an independently developed subsystem), obviously the combination of components or unit terms makes sense only when some compatibility conditions are fulfilled.

Let us now consider an example where the compatibility condition is violated.

³ In the application of the generic component G we need an explicit fitting symbol map since otherwise the sort Cont[Elem] can ambiguously be mapped to either Cont[Nat] or Set.

arch spec WRONG_ARCH_SPEC = units CN: CONT [NATURAL_ORDER]; SN: GENERATED_SET [NATURAL_ORDER]; F: CONT [ELEM] × GENERATED_SET [ELEM] \rightarrow CONT2SET [ELEM] result F [CN] [SN]

The architectural specification WRONG_ARCH_SPEC is a variant of ARCH_ CONT2SET_NAT where, instead of requiring a component N implementing NATURAL_ORDER and two generic components implementing containers and sets respectively, we just require a component CN implementing containers of natural numbers and a component SN implementing sets of natural numbers. However, then the application F [CN] [SN] makes no sense since there is no way to ensure that the common symbols of CN and SN have the same interpretation. It may indeed be the case that natural numbers are interpreted in some way in CN and in a different way in SN, which makes the application of F impossible. (Hence a similar problem would arise if one would use the combination of components 'CN and SN'.)

Let us now consider a more complex example.

arch spec BADLY_STRUCTURED_ARCH_SPEC = units N: NATURAL_ORDER; A: NATURAL_ORDER \rightarrow NATURAL_ARITHMETIC; C: CONT_COMP; S: SET_COMP; F: CONT [ELEM] \times GENERATED_SET [ELEM] \rightarrow CONT2SET [ELEM]

```
result F [C [A [N]]] [S [A [N]]]
```

The architectural specification BADLY_STRUCTURED_ARCH_SPEC is a variant of ARCH_CONT2SET_NAT where, in addition to the component N implementing NATURAL_ORDER, we require a generic component A which is used to expand N into an implementation of NATURAL_ARITHMETIC. In the architectural specification ARCH_CONT2SET_NAT, the compatibility condition in the application F [C[N]] [S[N]] was easy to discharge. Here, in the result unit term F [C[A[N]]] [S[A[N]]] of BADLY_STRUCTURED_ARCH_SPEC, we apply F to the pair made of C[A[N]] and S[A[N]]. In this case only a semantic analysis can ensure that these two arguments are compatible, since the common symbols cannot be traced to the same non-generic component, but only to two applications of the same generic component A to similar arguments. (Actually the arguments are just the same here, but in general checking this would require non-trivial semantic reasoning.)

It is advisable to use unit terms where compatibility can be checked by a simple static analysis. CASL provides additional constructs which make it easy to follow this recommendation, as explained below. Auxiliary unit definitions or local unit definitions may be used to avoid repetition of generic unit applications.

```
arch spec Well_Structured_Arch_Spec =
units N : NATURAL_ORDER:
         : NATURAL_ORDER \rightarrow NATURAL_ARITHMETIC;
      A
      AN = A[N];
      C : CONT_COMP;
      S : Set_Comp;
      F : Cont [ELEM] × GENERATED_SET [ELEM]
                   \rightarrow Cont2Set [ELEM]
result F[C[AN]][S[AN]]
arch spec ANOTHER_WELL_STRUCTURED_ARCH_SPEC =
units N: NATURAL_ORDER;
      A : NATURAL_ORDER \rightarrow NATURAL_ARITHMETIC;
      C : CONT_COMP;
      S : Set_Comp;
      F : Cont [Elem] × Generated_Set [Elem]
                  \rightarrow Cont2Set [ELEM]
```

result local AN = A[N] within F[C[AN]][S[AN]]

The problem illustrated in BADLY_STRUCTURED_ARCH_SPEC can be fixed easily. An auxiliary unit definition may be used to avoid the repetition of generic unit applications, such as (AN = A[N]) in Well_STRUCTURED_ ARCH_SPEC. An alternative is to make the definition of AN local to the result unit term, as illustrated in ANOTHER_WELL_STRUCTURED_ARCH_SPEC. In both cases common symbols can be traced to a non-generic unit, and compatibility can be checked by an easy static analysis.

Libraries

Libraries are named collections of named specifications.

In the foregoing chapters, we have seen many examples of named specifications, and of references to them in later specifications. This chapter explains how a collection of named specifications can itself be named, as a *library*. The creation of libraries facilitates the *reuse* of specifications. For practical applications, it is important to be able to reuse (at least) existing specifications of basic datatypes, such as those described in Chap. 12.

Local libraries are self-contained.

A library is called *local* when it is self-contained, i.e., for each reference to a specification name in the library, the library includes a specification with that name. Local libraries might appear at first sight to be all that we need, but actually they provide poor support for reuse of specifications. The problem is that when a specification from one local library is reused in another, it has to be repeated *verbatim*. There is no formal link between the original specification and the copy, despite them having the same name: the names used in a library can be chosen freely, and different libraries could use the same name for completely different specifications.

Distributed libraries support reuse.

Distributed libraries allow duplication of specifications to be avoided altogether. Instead of making an explicit copy of a named specification from one library for use in another, the second library merely indicates that the specification concerned can be *downloaded* from the first one.

Different versions of the same library are distinguished by hierarchical version numbers.

In practice, specifications *evolve*, e.g., to provide further operations or predicates on the specified sorts, or to define new subsorts. The libraries containing the specifications can evolve too, by adding or removing named specifications. Without some form of version control, even a trifling change in one library might cause specifications in other libraries to become ill-formed, or affect their meanings. CASL allows different versions of the same library to coexist (distinguishing them by hierarchical version numbers), and allows downloadings in a library to indicate that a particular version of another library is required.

Creation of new libraries is essential in connection with larger specification projects, and projects of any scale can benefit from reuse of specifications from existing libraries. The rest of this chapter illustrates the constructs used to specify local libraries, distributed libraries, and versions, and gives some advice on the organization of libraries.

9.1 Local Libraries

Local libraries are self-contained collections of specifications.

```
library UserManual/Examples
```

```
...

spec Natural = ...

spec Natural_Order = Natural then ...
```

The collection of all the illustrative examples given in the foregoing chapters is self-contained, so it could be made into a local library and named USER-MANUAL/EXAMPLES, as outlined above. To provide a separate local library for each chapter would however involve a considerable amount of duplication, since many of the specifications that are defined in the earlier chapters are also referenced in later chapters (e.g., SET_PARTIAL_CHOOSE in Chap. 4 instantiates GENERATED_SET, which is defined in Chap. 3). Using distributed libraries, as explained in Sect. 9.2, this duplication can be avoided.¹

The 'same name, same thing' principle of CASL applies only within specifications, and it is possible for a library to include alternative specifications for the same symbols (e.g., using different sets of axioms). However, when such alternative specifications are both extended (perhaps indirectly) in the same specification, the principle does apply, and unintended name clashes might then arise. Thus in general, it is advisable for the developers of a library to respect the 'same name, same thing' principle when choosing symbols throughout the library. In any case, this is obviously helpful to those who might later browse the library. Alternative specifications for the same symbols should therefore be given in separate libraries.²

Specifications can refer to previous items in the same library.

library UserManual/Examples

spec STRICT_PARTIAL_ORDER = ...

spec Total_Order = Strict_Partial_Order **then** ...

spec PARTIAL_ORDER = STRICT_PARTIAL_ORDER **then** ...

. . .

. .

Although we may often regard libraries as *sets* of named specifications, they are actually *sequences*, and the order in which the specifications occur is significant.

Specification names have linear visibility: each specification can refer only to the names of the specifications that precede it. Thus a series of extensions has to be presented in a bottom-up fashion, starting with a specification that is entirely self-contained, containing no references to other specifications at all. Each specification name in a library has a unique defining occurrence, so overriding cannot arise. Extensions that do not refer to each other may be given in any order (e.g., PARTIAL_ORDER above could just as well be given before TOTAL_ORDER).

Linear visibility of specification names means that mutual recursion between specifications is prohibited. When two specifications each make use of symbols declared in the other, the declarations of those symbols have to be duplicated, or moved to a preceding specification that can then be referenced by them both.

¹ A distributed library for each chapter of Part II is available via the CoFI web pages; copies are provided on the CD-ROM accompanying this book.

² If we intended our comprehensive USERMANUAL/EXAMPLES library for general use, we would remove all the illustrative alternative specifications.

All kinds of named specifications can be included in libraries.

```
library USERMANUAL/EXAMPLES
...
spec STRICT_PARTIAL_ORDER = ...
spec GENERIC_MONOID [sort Elem] = ...
...
view INTEGER_AS_TOTAL_ORDER : ...
...
view LIST_AS_MONOID [sort Elem] : ...
...
arch spec SYSTEM = ...
...
unit spec CONT_COMP = ...
...
```

Items in libraries can be any kind of named specification, as illustrated above: simple named specifications, generic specifications, named view definitions, generic view definitions, and architectural and unit specifications. We shall henceforth refer to them generally as *library items*.

Libraries themselves never include anonymous specifications, such as declarations of sorts and operations. Moreover, the symbols declared by a library item are *not* automatically available for use in subsequent items: an explicit reference to the name of the library item is required to 'import' the item.

Technically, each library item is said to be *closed*, being interpreted without any pre-declared symbols at all. This facilitates downloading items from distributed libraries, see Sect. 9.2.

Display, parsing, and literal syntax annotations apply to entire libraries.

library UserManual/Examples

. . .

```
%display __<=__ %LATEX _ ≤ _
%display __>=__ %LATEX _ ≥ _
%display __union__ %LATEX _ ∪ __
%prec {_+_, __-} < {_* _}
%left_assoc __+__, _ * _
...
spec Strict_PARTIAL_ORDER = ...
...
```

```
spec PARTIAL_ORDER = STRICT_PARTIAL_ORDER then ... \leq ... < ... < ... < ... < ... < ..
```

Annotations affecting the way terms are written or displayed apply to an entire library, and have to be collected at the beginning of the library. These annotations include display and precedence annotations, illustrated above.

Recall that various reserved words and symbols in CASL specifications are input in ASCII, but displayed as mathematical signs (e.g., universal quantification is input as 'forall', and displayed as ' \forall ' when this sign is available in the current display format). *Display annotations* provide analogous flexibility for declared symbols. For example, the display annotations illustrated above determine how infix symbols input as '<=', '>=', and 'union' are displayed when using LATEX to format the specification. Note that a display annotation applies to all occurrences of the input symbol in the library, regardless of overloading.

Display annotations can give alternative displays for different formats: apart from IATEX, both RTF and HTML are presently supported. The display of the annotation itself shows only the input syntax of the symbol and the result produced by the current formatter. The input form of one of the above annotations might be as follows:

%display __union__ %HTML __^U %LATEX __\cup__

When no display annotation is given for a particular format, the input format itself is displayed. Thus the symbol displayed as ' \cup ' in the present IAT_EX version of this User Manual would be displayed as '*union*' in an RTF version, unless the above annotation were to be extended with an RTF part.

Parsing annotations allow omission of grouping parentheses when terms are input. A single annotation can indicate the relative precedence or the associativity (left or right) of a group of operation symbols. The precedence annotation for infix arithmetic operations given above, namely:

 $\operatorname{mec} \{-+, --\} < \{-*\}$

allows a term such as a + (b * c) to be input (and hence also displayed) as a + b * c. The left-associativity annotation for + and *:

%left_assoc _+__, _ * _

allows (a + b) + c to be input as a + b + c, and similarly for *; but the parentheses cannot be omitted in (a + b) - c (not even if '____' were to be included in the same left-associativity annotation).

When an operation symbol is declared with the associativity attribute *assoc*, an associativity *annotation* for that symbol is provided automatically.³ Thus in practice, explicit associativity annotations are needed only for non-associative operations such as subtraction and division.

Libraries and library items can have author and date annotations.

An *author annotation* at the beginning of a library indicates the collective authorship of the entire library; one preceding an individual library item indicates its specific authorship.

A *date annotation* at the beginning of a library should indicate the release date of the current version of the library. It may also give the release dates of some previous major versions, possibly including that of the original version. A date annotation on an individual library item should indicate when that item was last changed, and (optionally) the dates of previous changes.

9.2 Distributed Libraries

Libraries can be installed on the Internet for remote access.

library Basic/Numbers

%left_assoc __@@__ %number __@@__ %floating _...._, __E__ %prec {__E__} < {_...._}

³ This implicit parsing annotation is local to the enclosing specification and to specifications that reference it, in contrast to ordinary parsing annotations.

```
spec NAT =

free type Nat ::= 0 | suc(Nat)

...

ops 1 : Nat = suc(0); ...; 9 : Nat = suc(8);

...

spec INT = NAT then ...

spec RAT = INT then ...

spec DECIMALFRACTION = RAT then

...

ops __:::__: Nat × Nat \rightarrow Rat;

...

\_E\_: Rat \times Int \rightarrow Rat
```

The above example is an extract from one of the CASL libraries of basic datatypes, described in Chap. 12 and available on the Internet. It illustrates the overall structure of a library intended for general use, as well as some helpful annotations concerning literal syntax for numbers, which are explained below.

Validated libraries can be registered for public access.

CoFI will maintain a register of useful libraries. Registered CASL libraries are identified by hierarchical path names. For instance, all the CASL libraries of basic datatypes have names starting with 'BASIC/', and path names starting with 'CASL/' are reserved for libraries connected with the CASL language itself (e.g., the specification of the abstract syntax of CASL in CASL).

Registered libraries will be mirrored at several sites, to ensure their continuous accessibility. The URLs of a library can be obtained from the library name using a table provided on the CoFI web pages.

Libraries have to be *validated* before registration. The validation of a library ensures not only that it is well-formed, but also that semantic annotations expressing consistency of specifications (or conservativity over the parameters, in case of generic or unit specifications) have been added, and that all proof obligations (corresponding both to well-formedness conditions and to semantic annotations in the library) have been verified.

It is likely that new versions of existing libraries will be produced, e.g., providing further operations whose usefulness was not realized beforehand. Although the assignment and use of library version numbers allows users to protect their specifications from changes due to new versions (see Sect. 9.3), at least the names used in a registered library should not change much between versions.

Libraries should include appropriate annotations.

In particular, parsing and display annotations can be provided, as explained in Sect. 9.1. The above example illustrates a further kind of annotation, used to provide *literal syntax* for numbers in CASL. The effect of the illustrated annotations is that, after downloading the appropriate specifications from the library BASIC/NUMBERS, conventional decimal notation can be used for integers and decimal fractions, e.g., 42, 2.718, 10E - 12. The digits 42 are interpreted as the term 4@@2, and 2.718 is interpreted as the term 2:::718 (where 718 is subsequently interpreted as (7@@1)@@8). The definition of the operation $_@@_$ is shown above; those of $_:::_$ and $_E_$ are a bit more involved, and omitted here. Notice that the library BASIC/NUMBERS is *not* hard-wired into CASL, and users could provide annotations to interpret the literal syntax for integers and decimal fractions as terms involving different operations, e.g., on different sorts.

Libraries can include items downloaded from other libraries.

library BASIC/STRUCTUREDDATATYPES

```
from BASIC/NUMBERS get NAT, INT
...
spec LIST [sort Elem] given NAT = ...
...
spec ARRAY ... given INT = ...
```

Individual specifications and other items can be *downloaded* from other libraries. For example, the library BASIC/STRUCTUREDDATATYPES, outlined above, does not itself provide the specifications of natural numbers and integers that are needed in the specifications LIST and ARRAY, but instead downloads NAT and INT from the BASIC/NUMBERS library.

The names of the items to be downloaded from a library have to be listed explicitly: one cannot request the downloading of all the items that happen to be provided by a library. However, although the name of each item provided by a downloading is always explicit, no indication is given of its kind (i.e., whether it is an ordinary, generic, or architectural specification, or a view) nor of what symbols it declares. Thus the well-formedness of a library depends on what items are actually downloaded from other libraries.

The construct '**from** ... **get** ... ' above has the effect of downloading the specifications that are named NAT and INT in BASIC/NUMBERS, preserving their names. It is also possible to give downloaded specifications different

names, e.g., to avoid clashes with specification names that are already in use locally:

from Basic/Numbers get Nat \mapsto Natural, Int \mapsto Integer

Items that are referenced by downloaded items are *not* themselves automatically downloaded, e.g., downloading INT does not entail the downloading of NAT. This is because downloading involves the *semantics* of the named items, not their *text*. The semantics of INT consists of a signature and a class of models, and is a self-contained entity – recall from Chap. 6 that the models of a structured specification have no more structure than do those of a flat, unstructured specification. Thus downloading INT gives exactly the same result as if the reference to NAT in its text had already been replaced by the text of NAT *before* downloading. For the same reason, the presence of another item with the name NAT in the current library makes no difference to the result of downloading INT. In terms of software packages, downloading specifications from CASL libraries is analogous to installing packages from binaries, rather than from sources.

Downloading any item from another library B in a library A causes all the parsing and display annotations of B to be inserted at the beginning of A. (Conflicting annotations from different libraries are ignored altogether, and local annotations override conflicting downloaded annotations.) The copied annotations allow terms to be written and displayed in A in the same way as in B.

Substantial libraries of basic datatypes are already available.

The organization of the following libraries of basic datatypes is explained in Chap. 12:

BASIC/NUMBERS: natural numbers, integers, and rationals.

BASIC/RELATIONSANDORDERS: reflexive, symmetric, and transitive relations, equivalence relations, partial and total orders, boolean algebras.

BASIC/ALGEBRA_I: monoids, groups, rings, integral domains, and fields. BASIC/SIMPLEDATATYPES: booleans, characters.

 ${\tt BASIC/StructuredDatatypes:\ sets, lists, strings, maps, bags, arrays, trees.}$

BASIC/GRAPHS: directed graphs, paths, reachability, connectedness, colorability, and planarity.

BASIC/ALGEBRA_II: monoid and group actions on a space, euclidean and factorial rings, polynomials, free monoids, and free commutative monoids.

BASIC/LINEARALGEBRA_I: vector spaces, bases, and matrices.

BASIC/LINEARALGEBRA_II: algebras over a field.

BASIC/MACHINENUMBERS: bounded subtypes of naturals and integers.

These libraries form a coherent collection of highly-polished specifications. The libraries themselves are provided in full in the CASL *Reference Manual* [20], and are available on the Internet.

Libraries need not be registered for public access.

library http://www.cofi.info/CASL/Test/Security

```
from http://casl:password@www.cofi.info/CASL/RSA get KEY
...
spec DECRYPT = KEY then ...
...
```

Libraries under development, and libraries provided for restricted groups of users, are named and accessed by their URLs (instead of the simple path names used for registered libraries). This allows the CASL library constructs to be fully exploited for libraries that are not yet – and perhaps never will be – registered for public access. Moreover, validation of libraries can be a demanding and time-consuming process, and getting a library approved and registered is appropriate only when it provides specifications that are likely to be found useful by persons not directly involved in its development.

The primary Internet access protocols HTTP and FTP both support password protection of libraries and the insertion of usernames and passwords in URLs allows downloading between protected libraries (With HTTP, the username and password can be unrelated to those used for the host file system.)

9.3 Version Control

Subsequent versions of a library are distinguished by explicit version numbers.

library Basic/Numbers version 1.0

```
spec NAT = ...

spec INT = NAT then ...

spec RAT = INT then ...

...
```

As illustrated above, a library can be assigned an explicit version number, allowing it to be distinguished from previous and future versions of the same library. CASL allows conventional hierarchical version numbers, familiar from version numbers of software packages: the initial digits indicate a major version, digits after a dot indicate sub-versions, and digits after a further dot indicate patches to correct bugs. (Distinctions between alpha, beta, and other pre-release versions are not supported.)

The smallest version number is written simply '0', and can be omitted when specifying the initial version of a library; this is the case with the version of BASIC/NUMBERS shown in Sect. 9.2, it implicitly has version number '0', but in general the first-installed version of a library could have any version number at all. The numbers of successively installed versions do not have to be contiguous, nor even increasing: e.g., a patched version 0.99.1 could be installed after version 1.0.

Individual library items do not have separate version numbers. Date annotations can be used to indicate which items have changed between two versions of a library.

Libraries can refer to specific versions of other libraries.

```
library BASIC/STRUCTUREDDATATYPES version 1.0
```

```
from BASIC/NUMBERS version 1.0 get NAT, INT
...
spec LIST [sort Elem] given NAT = ...
...
spec ARRAY ... given INT = ...
```

Downloading items from particular versions of libraries is necessary if one wants to ensure coherence between libraries. For example, as illustrated above, version 1.0 of BASIC/STRUCTUREDDATATYPES downloads NAT and INT from version 1.0 of BASIC/NUMBERS. Omitting the version number when downloading gives implicitly the *current version* of the library, which may of course change. By the way, the current version of a library is *not* necessarily the one most recently installed: it is the one with the *largest version number*. As previously mentioned, a patched version 0.99.1 could be installed after version 1.0, but a downloading without an explicit version number would still refer to version 1.0.

Even though the developers of libraries may try to ensure backwards compatibility between versions, it could happen that symbols introduced in a new version of a downloaded specification clash with symbols already in use in the library that specified the downloading, causing ill-formedness or inconsistency. So for safety, it is advisable to give explicit version numbers when downloading (also when downloading from version '0' of another library). If one subsequently wants to use symbols that are introduced only in some later version of another library, all that is needed is to change the version number in the downloading(s).

An alternative strategy is to ensure consistency with the *current* versions of all libraries from which specifications are downloaded, by observing the changes in the new versions and adapting the downloading library accordingly. For instance, one might download INT from the current version of BA-SIC/NUMBERS, instead of from version 1.0 of that library. This may involve extra work when a new version of BASIC/NUMBERS appears, but it has several advantages over the more cautious approach. CASL leaves the choice to the user, although registered libraries will generally be required to use explicit version numbers when downloading from other libraries.

All downloadings should be collected at the beginning of a library.

Although CASL allows downloadings to be interleaved with specification definitions, it is advisable to collect the downloadings at the beginning of libraries (together with any parsing and display annotations). This makes it easy to see dependencies between libraries, and to ensure that different downloadings from the same library all refer to the same version of it.

Foundations

Donald Sannella and Andrzej Tarlecki

A complete presentation of CASL is in the Reference Manual.

This User Manual has introduced the potential user to the features of CASL mainly by means of illustrative examples. It has presented and discussed the typical ways in which the language concepts and constructs are expected to be used in the course of building system specifications. Thus, the presentation in Part II focused on *what* the constructs and concepts of CASL are *for*, and *how* they should (and should not) be used. We tried to make these points as clear as possible by referring to simple examples, and by discussing both the general ideas and some details of CASL specifications. We hope that this has given the reader a sufficient feel of the formalism, and enough understanding, to look through the presentation of a basic library of CASL specifications in Chap. 12, to follow the case study in Chap. 13, and to start experimenting with the use of CASL for writing specifications – perhaps employing the support tools presented in Chap. 11.

By no means, however, should this book be regarded as a complete presentation of the CASL specification formalism – this is given in the accompanying volume, the CASL *Reference Manual* [20].

CASL has a definitive summary.

The CASL Reference Manual begins with a definitive *summary* of the entire CASL language: all the language constructs are listed there systematically, together with the syntax used to write them down and a detailed explanation of their intended meaning. However, although it tries to be precise and complete, the CASL Summary still relies on natural language to present CASL. This inherently leaves some room for interpretation and ambiguity in various corners of the language, for example where details of different constructs interact.

CASL has a complete formal definition.

A key aim of the entire CoFI initiative is to avoid any potential ambiguities by providing a complete formal definition for CASL, and sound mathematical foundations for the advocated methodology of its use in software specification and development.

Abstract and concrete syntax of CASL are defined formally.

The next part of the Reference Manual is a formal definition of the syntax of CASL. *Abstract syntax* is given, where each phrase is written in a way that directly indicates its components, thus making evident its internal structure. In essence, the use of each construct of the language is explicitly labeled here. This is convenient for formal manipulation and analysis, but is not so readable. Therefore, the so-called *concrete syntax* of CASL (as used for instance in the examples throughout this book) is given as well, retaining a direct correspondence with the abstract syntax. This offers to the user of CASL a convenient and readable way of writing down CASL specifications, in a way that makes clear the formal structure of the phrases and constructs used to build them. As usual, the syntax is given as a context-free grammar, using a variant of the BNF notation, relying on well-established theory to give its formal meaning, and on a variety of tools and techniques available for syntactic analysis of languages presented in such a style.

CASL has a complete formal semantics.

The ultimate definition of the meaning of CASL specifications is provided by the *semantics* of CASL in the Reference Manual. The semantics first defines mathematical entities that formally model the intended meaning of various concepts underlying CASL, as already hinted at in Chap. 2, and further introduced and discussed throughout the summary.

The key concepts here are that of CASL signature, model and formula, together with the satisfaction relation between models and formulas over a common signature. In fact, these are variants of the standard algebraic and logical notions, thus linking work on CASL to well-established mathematical theories and ideas.

In a more or less standard way, we use these concepts to build the *semantic* domains which the meanings of phrases in various syntactic categories of CASL inhabit. We have chosen to give the semantics in the form of so-called *natural* semantics, with formal deduction rules to derive judgments concerning the meaning of each CASL phrase from the meanings of its constituent parts. Not only is this a mathematically well-established formalism with an unambiguous interpretation, but we also hope that this makes the semantics itself more readable, with details easier to follow than in some other approaches.

The overall semantics of CASL consists of two parts. The *static semantics* captures a form of static analysis of CASL specifications, in which they are checked for well-formedness – for example, that axioms are well-typed, and references to named entities are in scope. Then the *model semantics* takes a well-formed CASL specification and assigns a model-theoretic meaning to it.

CASL specifications denote classes of models.

In CASL, well-formed specifications denote signatures (static semantics) and classes of models (model semantics). Basic specifications, which in essence present a signature and a set of axioms over this signature, denote the class of models that satisfy all the axioms. The semantics of basic specifications is split into two parts: first many-sorted basic specifications are described and then features for defining and using subsorts are added.

The semantics is largely institution-independent.

A few more concepts are needed to explain the semantics of structured specifications. Key here is the notion of signature morphism, and the model reducts and translation of sentences that signature morphisms induce. Having introduced those, we obtain the *institution* [23] of CASL, that is, the underlying logical system of CASL equipped with extra structure to capture ways of moving between signatures linked by signature morphisms. One important point of the semantics is that all the layers of the semantics "above" basic specifications are *institution-independent*, i.e., well-defined for any institution chosen to build basic specifications (as long as the institution comes with a bit of extra structure concerned with forming unions of signatures and defining signature morphisms).

Next, we have the semantics of architectural specifications, which relies on a formal counterpart of the concept of a unit (module) of a system to be developed: self-contained units are viewed simply as models of the underlying institution, and parametrized units as functions mapping such parameter models to result models. Architectural specifications provide a way of specifying the component units of a system and indicating how the overall system is built by putting these units together. This intuition is captured by the semantics of architectural specifications, which denote a class of permitted unit bindings and a function that maps each such environment to a result unit.

Finally, the libraries layer of CASL is given a rather standard description with libraries modeled as environments giving names to the entities introduced (specifications, architectural specifications, etc.).

The semantics is the ultimate reference for the meanings of all CASL constructs.

Overall, the formal mathematical semantics is crucial in the understanding of CASL specifications. It provides their unambiguous meaning, and thus gives the ultimate reference point for all questions concerning the interpretation of any CASL phrase in any context.

We have already experienced how important such a formal semantics may be in the design of CASL itself. In many cases, the intended semantics was prominent in internal discussions on the details of the constructs under consideration, and provided guidelines for many choices in the design of CASL. Indeed, the concrete syntax of CASL was designed only after the semantics was settled.

Proof systems for various layers of CASL are provided.

The semantics is also a necessary prerequisite for the development of mechanisms for formal reasoning about CASL specifications. This is the role of *proof calculi* that support reasoning about the various layers of CASL. The starting point is a formal system of deduction rules which determines a proof-theoretic counterpart of the consequence relation between sets of formulas, thus providing a way for deriving consequences of sets of axioms in CASL specifications. This is then extended to systems of rules for deriving consequences of structured specifications and for proving inclusions between classes of models of such specifications. These systems are also used in rules for formal verification of the internal correctness of system designs as captured by architectural specifications. For all these systems, their soundness is proved and completeness discussed by reference to the formal semantics of CASL.

One point of interest is that, again, the extension of the basic proof system for consequences between sets of formulas to structured and architectural specifications does not really rely on the specifics of the underlying institution, but just reflects the way in which the structuring and architectural constructs are defined for an arbitrary institution. The foundations of our CASL are rock-solid!

All this work on the mathematical underpinnings of CASL and related specification and development methodology, as documented in the Reference Manual, should make it exceptionally trustworthy – at least in the sense that it provides a formal point of reference against which all the claims may (and should) be checked.

Tools

Till Mossakowski

This chapter gives an overview of the CASL tools. Analysis tools for CASL like parsers and static checkers, as well as formatters, are stable now and cover the whole of CASL. Proof tools are available but are less mature.

CASL has been designed with the goal of subsuming many previous specification languages. Most of these languages come with specific tools, and of course, these tools should be reusable in the context of CASL. Hence, a central issue is to build bridges to existing tools (rather than building new tools from scratch). Using an interchange format generated by the analysis tools, CASL has been interfaced in this way to rewriting engines and theorem provers, usually working for a subset of CASL.

Naturally, due to the ongoing development of these tools, detailed descriptions would become outdated sooner or later. Therefore, we give here just an appetizer, intended to encourage the reader to install the tools and experiment with them (and to convince her/him that this is rather easy). More detailed descriptions of the tools, as well as their latest versions and other tools that may be developed in the future, are available by following the links on the CoFI tools home page [21]: http://www.cofi.info/Tools.

The analysis tools for CASL have been used to check all the examples contained in this book, as well as the CASL Basic Libraries [20]. Moreover, some proofs from a case study in refinement have been carried out with the proof tools.

CASL specifications can be checked for well-formedness using a form-based web page.

The easiest way to check a CASL specification for well-formedness is to visit the web interface. Using a web form, you can submit your specification (without the need to install anything on your machine), and get immediate feedback about the well-formedness of the specification. Follow the "web interface" link on the CoFI tools home page.

11.1 The Heterogeneous Tool Set (HETS)

The Heterogeneous Tool Set (HETS) is the main analysis tool for CASL.

HETS is a tool set for the analysis of specifications written in CASL, its extensions and sublanguages – hence the name "heterogeneous". HETS consists of logic-specific tools for the parsing and static analysis of the different CASL extensions and sublanguages, as well as a logic-independent parsing and static analysis tool for structured and architectural specifications and libraries (which of course calls the logic-specific tools whenever a basic specification is encountered). In order to be able to tackle proof obligations occurring in (statically well-formed) specifications, HETS is interfaced with various logicspecific theorem proving, rewriting and consistency checking tools. On top of this, there is a logic-independent proof engine called MAYA, which manages the proof obligations. MAYA uses so-called *development graphs*, a graphical representation of CASL structured specifications.

The architecture of HETS is shown in Fig. 11.1. The latest version can be obtained from the CoFI tools home page [21]. Installation is easy; just follow the instructions.

Consider the first example in this book:

```
spec STRICT_PARTIAL_ORDER =

sort Elem

pred ... < ...: Elem × Elem

\forall x, y, z : Elem

• \neg(x < x) %(strict)%

• x < y \Rightarrow \neg(y < x) %(asymmetric)%

• x < y \land y < z \Rightarrow x < z %(transitive)%

%{ Note that there may exist x, y such that

neither x < y nor y < x. }%
```

```
end
```

Hets can be used for parsing and checking static well-formedness of specifications.



Fig. 11.1. Architecture of the heterogeneous tool set.

Let us assume that the example is in a file named Order.casl (actually, this file is provided on the web and on the CD-ROM coming with this volume). Then you can check the well-formedness of the specification by typing (into some shell):

hets Order.casl

HETS checks both the correctness of this specification with respect to the CASL syntax, as well as its correctness with respect to the static semantics (e.g. whether all identifiers have been declared before they are used, whether operators are applied to arguments of the correct sorts, whether the use of overloaded symbols is unambiguous, and so on).

It is also possible to generate a pretty printed LATEX version of Order.casl by typing:

hets -o pp.tex Order.casl

One use of Order.casl might be to express the fact that the natural numbers form a strict partial order as a view, as follows:

```
spec NATURAL = free type Nat ::= 0 \mid suc(Nat) end
```

```
spec NATURAL_ORDER_2 =
NATURAL

then pred __< _: Nat × Nat
\forall x, y : Nat
• 0 < suc(x)
• \neg x < 0
• suc(x) < suc(y) \Leftrightarrow x < y

end

view V1 : STRICT_PARTIAL_ORDER to NATURAL_ORDER_2 =
Elem \mapsto Nat
```

end

Again, these specifications can be checked with HETS. However, this only checks syntactic and static semantic well-formedness – it is *not* checked whether the predicate '___ < __' introduced in NATURAL_ORDER_2 actually is constrained to be interpreted by a strict partial ordering relation. Checking this requires theorem proving, which will be discussed below.

HETS also displays and manages proof obligations, using development graphs.

However, before coming to theorem proving, let us first inspect the proof obligations arising from a specification. This can be done with:

```
hets -g Order.casl
```

(assuming that the above two specifications and the view have been added to the file Order.cas1). HETS now displays a so-called development graph (which is just an overview graph showing the overall structure of the specifications in the library), see Fig. 11.2.



Fig. 11.2. Sample development graph.

Nodes in a development graph correspond to CASL specifications. Arrows show how specifications are related by the structuring constructs.

The solid arrow denotes an ordinary import of specifications (generated by the extension), while the dashed¹ arrow denotes a proof obligation (corresponding to the view). This proof obligation needs to be discharged in order to show that the view is well-formed.

As a more complex example, consider the following loose specification of a sorting function, taken from Chap. 6:

```
spec LIST_ORDER_SORTED

[TOTAL_ORDER with sort Elem, pred _ < __] =

LIST_SELECTORS [sort Elem]

then local pred __is_sorted : List

\forall e, e' : Elem; L : List

• empty is_sorted

• cons(e, empty) is_sorted
```

• cons(e, cons(e', L)) is_sorted \Leftrightarrow (cons(e', L) is_sorted $\land \neg(e' < e))$

within op $order : List \rightarrow List$

 $\forall L : List \bullet order(L) is_sorted$

\mathbf{end}

The following specification of sorting by insertion also is taken from Chap. 6:

```
spec LIST_ORDER [TOTAL_ORDER with sort Elem, pred ___ < __ ] = LIST_SELECTORS [sort Elem]
```

- then local op $insert : Elem \times List \rightarrow List$ $\forall e, e' : Elem; L : List$
 - insert(e, empty) = cons(e, empty)
 - insert(e, cons(e', L)) = cons(e', insert(e, L)) when e' < e

```
else \ cons(e, cons(e', L))
```

within op $order : List \rightarrow List$

- $\forall e: Elem; \ L: List$
- order(empty) = empty
- order(cons(e, L)) = insert(e, order(L))

end

Both specifications are related. To see this, we first inspect their signatures. This is possible with:

hets -g Sorting.casl

¹ Actually, the dashed arrow will be displayed as solid and in red by HETS; we do not have colors available here.

assuming that Sorting.casl contains the above specifications. HETS now displays a more complex development graph, see Fig. 11.3.



Fig. 11.3. Development graph for the two sorting specifications.

Internal nodes in a development graph correspond to unnamed parts of a structured specification.

In the above-mentioned development graph, we have two types of nodes. The named ones correspond to named specifications, but there are also unnamed nodes corresponding to anonymous basic specifications like the declaration of the *insert* operation in LIST_ORDER above. Basically, there is an internal node for each structured specification that is not named.

Again, the simple solid arrows denote an ordinary import of specifications (corresponding to the extensions and unions in the specifications), while the double arrows denote hiding (corresponding to the local specification).

By clicking on the nodes, one can inspect their signatures. In this way, we can see that both LIST_ORDER_SORTED and LIST_ORDER have the same signature. Hence, it is legal to add a view:

view v2[Total_Order] : List_Order_Sorted[Total_Order] to List_ Order[Total_Order]

end

We have already added this view to Sorting.casl. The corresponding proof obligation between LIST_ORDER_SORTED and LIST_ORDER is displayed in Fig. 11.3 as a dotted arrow.

Proof obligations can be discharged in various ways.

Trivial proof obligations can be discharged by HETS alone using the "Proofs" menu. The proof obligation in Fig. 11.3, indicated by the lower dotted arrow between LIST_ORDER_SORTED and LIST_ORDER, states that insertion sort, as defined by the operation *order* in LIST_ORDER, actually has the properties of a sorting algorithm. Here, one has to choose a theorem prover that is to be used to discharge the proof obligation, which is then done by using commands specific to the theorem prover (cf. e.g. Section 11.2). Alternatively, one can state that one just conjectures the obligation to be true.

HETS is written in Haskell. Its parser uses combinator parsing. The userdefined (also known as "mixfix") syntax of CASL calls for a two-pass approach. In the first pass, the skeleton of a CASL abstract syntax tree is derived, in order to extract user-defined syntax rules. In a second pass, which is performed during static analysis, these syntax rules are used to parse any expressions that use mixfix notation. The output is stored in the so-called ATerm format [12], which is used as interchange format for interfacing with other tools.

HETS provides an abstract interface for institutions, so that new logics can be integrated smoothly. In order to do so, a parser, a static checker and a prover for basic specifications in the logic have to be provided. HETS has been built based on experiences with its precursors, CATS and MAYA. The CASL Tool Set (CATS) [28, 30] comes with roughly the same analysis tools as HETS. The management of development graphs is not integrated in CATS, but is provided with a stand-alone version of the tool MAYA [5, 4]. CATS and MAYA can be obtained from the CoFI tools home page [21].

11.2 Hol-Casl

Hol-Casl is an interactive theorem prover for Casl, based on the tactical theorem prover Isabelle.

The HoL-CASL system [28] provides an interface between CASL and the theorem proving system ISABELLE/HoL [31]. We have chosen ISABELLE because it has a very small core guaranteeing correctness, and its provers, like the simplifier or the tableaux prover, are built on top of this core. Furthermore, there is over ten years of experience with it, and several mathematical textbooks have been partially verified with ISABELLE.

CASL is linked to ISABELLE/HOL by an encoding.

Since subsorting and partiality are present in CASL but not in ISABELLE/HOL, we have to encode these features, as explained in [29]. This means that theorem proving is not done in the CASL logic directly, but in the logic HOL (for higher-order logic) of ISABELLE. HOL-CASL tries to make the user's life easy by:

- choosing a shallow embedding of CASL into HOL, which means that e.g. CASL's logical implication => is mapped directly to ISABELLE/HOL's logical implication --> (and the same holds for other logical connectives and quantifiers); and
- adapting ISABELLE/HOL's syntax to conform with the CASL syntax, e.g. ISABELLE/HOL's --> is displayed as =>, as in CASL.

However, it is essential to be aware of the fact that the ISABELLE/HOL logic is different from the CASL logic. Therefore, the formulas appearing in subgoals of proofs with HOL-CASL will not fully conform to the CASL syntax: they may use features of ISABELLE/HOL such as higher-order functions that are not present in CASL. HOL-CASL can be obtained from the CoFI tools home page [21].

To start the HoL-CASL system, follow the installation instructions, and then type:

HOL-CASL

You can load the above specification file Order.casl by typing:

use_casl "Order";

Let us try to prove part of the view V1 above. To prepare for conducting a proof in the target specification of the view, NATURAL_ORDER_2, type in:

CASL_context Natural_Order_2.casl;

AddsimpAll();

The first command just selects the specification as the current proof context; the second one adds all the axioms of the specification to ISABELLE's simplifier (a rewriting engine). Note that the latter is advisable only if the axioms are terminating, when considered as a set of rewrite rules.

To prove the first property expressed by the view, we first have to type in the goal. Then we chose to perform induction over the variable x, and the rest can be done with automatic simplification. Finally, we name the theorem for later reference:

```
Goal "forall x:Nat . not x<x";
by (induct_tac "x" 1);
by Auto_tac;
ged "Nat_irreflexive";
```

Both the stand-alone MAYA as well as the MAYA part of HETS also provide an interface to HOL-CASL, so that it can be used to discharge proof obligations arising in development graphs.

11.3 AsF+SDF Parser and Syntax-Directed Editor

ASF+SDF was used to prototype the CASL syntax.

The algebraic specification formalism AsF+SDF [22] and the AsF+SDF Meta-Environment [13] have been deployed to prototype CASL's concrete syntax, and to develop a mapping for the concrete syntax to an abstract syntax representation using so-called ATerms [12]. Currently, only the first pass of parsing (i.e. without mixfix analysis) is realized in SDF. Parsing is performed based on the underlying Scannerless Generalized LR parsing technology. A prototype of the mapping from the concrete to abstract representation is written in ASF rewrite rules. The Asf+SDF Meta-Environment provides syntax-directed editing of CASL specifications.

Given the concrete syntax definition of CASL in SDF, syntax-directed editors within the ASF+SDF Meta-Environment come for free. Recent developments in the Meta-Environment even allow for the development of a CASL specification environment.

The ASF+SDF Meta-Environment provides a built-in library mechanism which contains a collection of grammars, among others CASL. Via the library the user of the Meta-Environment has access to the CASL syntax in SDF and a collection of ASF equations to map CASL specifications into an interchange format named CasFix. The asfc tool compiles the CASL grammar into a stand-alone C program.

A link to the AsF+SDF Meta-Environment with further information and a download possibility can be found at the CoFI tools home page [21]. The built-in library of the Meta-Environment (Version 1.5 and higher) will provide the full CASL grammar in SDF and the mapping to CasFix as an ASF specification.

11.4 Other Tools

The following tools are at a prototype stage at the time of appearance of this volume. Please refer to at the CoFI tools home page [21], where the latest versions can be downloaded.

Translation to OCAML

A translation from CASL into OCAML has been developed at Paris and Poitiers. The translation works for a "functional programming" sublanguage of CASL that includes free datatypes and recursive definitions of operations over these types.

Translation to Haskell

A translation from CASL into Haskell has been developed in Bremen. Actually, the translation works on an executable subset of HASCASL (a higher-order extension of CASL). Using an embedding of CASL into HASCASL, this translation can also be used for CASL.

Elan-Casl

ELAN-CASL is a rewriting engine for the HORN⁼ sublanguage of CASL. It is based on a translation of that sublanguage to the input language of the rewriting engine ELAN.

Given a CASL specification and a query term, ELAN-CASL computes the normal forms of the term. Note that because the set of rules is not required to be terminating nor confluent, a query term may have several normal forms, or may not terminate.

ELAN can be called either as an interpreter or as a compiler.

CASL consistency checker

The CASL consistency checker (CCC) has been developed in Bremen and Swansea. With CCC, one can interactively check whether a CASL specification is consistent. It is a faithful implementation of a consistency calculus for full CASL [32].

Besides using certain syntactical criteria, the consistency calculus relies heavily on the CASL structuring mechanisms and their semantic annotations. Consequently, consistency proofs follow the structure of the given specification. In this way, the calculus highlights the (usually few) 'hot spots' of a specification (e.g. views requiring theorem proving), while the (lengthy) 'trivial' parts of the consistency argument are discharged automatically. As the consistency calculus works along the structure of the specification, the need to construct (and prove correct) actual models of specifications is avoided as far as possible.



Fig. 11.4. CCC System Architecture

The CASL Consistency Checker consists of four parts (see Fig. 11.4): first, a *logical core* implementing a representation of the propositions to be proved, and the basic proof rules, along with a generic unification package for the CASL abstract syntax; secondly, representations of the *calculi*: the base rules which are stated axiomatically, and derived rules; thirdly, *proof procedures* (automatic or semi-automatic decision procedures), which serve to discharge specific proof obligations; and finally, *proof infrastructure* such as a package which helps users to conduct goal-directed proofs, tactics that support writing advanced proof procedures, and a simple database which allows storage and retrieval of proved theorems. The system is encapsulated by a single interface which does not allow the end-user (i.e. the working specifier) to add any more axiomatic rules or provers; thus, the typing is used to both increase confidence in the correctness of the implementation in the style of LCF, and to protect the integrity of the system from user intervention.

Basic Libraries

Till Mossakowski

The CASL Basic Libraries contain the standard datatypes.

The CASL Basic Libraries consist of specifications of often-needed datatypes and views between them, freeing the specifier from re-inventing well-known things. This can be compared to standard libraries in programming languages. While this book often discusses several styles of specification with CASL, the basic datatypes consistently follow a specific style described in [20].

Here we show a cut-down version without axioms.

Here, we describe two of the libraries (see the overview in Fig. 12.1): the libraries of numbers and of structured datatypes. We also provide strippeddown versions of the libraries themselves, with some of the specifications and all axioms and annotations removed. These stripped-down versions can serve for getting a first overview of the signatures of the specified datatypes.

The full CASL Basic Libraries with complete specifications is presented in the CASL Reference Manual [20], and is also included in the CD-ROM coming with this volume. The latest version is available at:

http://www.cofi.info/Libraries

Hets can be used to get an overview of the Basic Libraries.

The HETS tool described in Chap. 11 allows the structure of the specifications in the libraries to be displayed as a graph and their signatures to be inspected. This is recommended as a way of obtaining a better overview, and also for answering specific questions that arise when using the basic datatypes.


Fig. 12.1. Dependency graph of the libraries of basic datatypes.

12.1 Library BASIC/NUMBERS

This library provides specifications of natural numbers, integers, and rational numbers.

The natural numbers are specified as a free datatype.

In the specification NAT, the natural numbers are specified as a free datatype, together with a collection of predicates and operations over the sort *Nat* of natural numbers.

Note that the names for the *partial operations* subtraction $_-?_$ and division $_/?_$ include a question mark. This is to avoid overloading with the *total operations* $_ _-$ on integers and $_-/_$ on rationals, which would lead to inconsistencies as both these specifications import the specification NAT. The *total operation* for subtraction differs from subtraction on the integers as

well. It is written $_-!_$, and it is 0 whenever the partial subtraction -? is undefined, while it otherwise coincides with the latter.

The digits are introduced as constants, together with an operation <u>__@@__</u> for concatenation of digits. Together with an annotation (see Chap. 9):

```
%number _@@_
```

this allows one to write the usual literals (like e.g. 8364) for natural numbers.

The introduction of the subsort *Pos*, consisting of the positive naturals, gives rise to certain new operations, e.g.:

 $_\times_: Pos \times Pos \rightarrow Pos$,

whose semantics is completely determined by overloading.

library BASIC/NUMBERS

spec NAT =free type $Nat ::= 0 \mid suc(pre:?Nat)$ preds $_\leq_, _\geq_, _<_, _>_: Nat \times Nat;$ even, odd : Natops $_!: Nat \rightarrow Nat;$ $_+_, _\times_, _^, min, max, _-!_: Nat \times Nat \rightarrow Nat;$ $_-?_, _/?_, _div_, _mod_, gcd : Nat \times Nat \rightarrow? Nat$ %% Operations to represent natural numbers with digits: 1: $Nat = suc(\theta);$ %(1_def_Nat)% ops 2: Nat = suc(1);%(2_def_Nat)% 3: Nat = suc(2);%(3_def_Nat)% 4: Nat = suc(3);%(4_def_Nat)% 5: Nat = suc(4);%(5_def_Nat)% 6: Nat = suc(5);%(6_def_Nat)% 7: Nat = suc(6);%(7_def_Nat)% 8: Nat = suc(7);%(8_def_Nat)% 9: Nat = suc(8); %(9_def_Nat)% $_@@_(m: Nat; n: Nat): Nat = (m \times suc(9)) + n$ %(decimal_def)% sort $Pos = \{p: Nat \bullet p > 0\}$ ops 1: $Pos = suc(\theta);$ $\%(1_as_Pos_def)\%$ $_\times_: Pos \times Pos \rightarrow Pos;$ $_+_: Pos \times Nat \rightarrow Pos;$ $_+_: Nat \times Pos \rightarrow Pos;$ $suc: Nat \rightarrow Pos$

end

The integers are specified as difference pairs of naturals.

The specification INT of integers is built on top of the specification NAT: integers are defined as equivalence classes of pairs of naturals written as differences — the axioms (which are omitted in the specification below) specify that two pairs are equivalent if their differences are equal:

$$\forall a, b, c, d: Nat$$

• $a - b = c - d \Leftrightarrow a + d = c + b$ %(equality_Int)%

The sort Nat is then declared to be a subsort of Int. Besides the division operator $__/?_$, the specification INT also provides the function pairs div/mod and quot/rem, respectively, as constructs for division — behaving differently on negative numbers, see [20] for a discussion. The operation sign gives the sign of an integer (which is either -1, 0, or 1).

```
spec INT =
       NAT
then generated type Int ::= \_-\_(Nat; Nat)
               Nat < Int
       sort
               \%\% a system of representatives for sort Int is
               \%\% a - 0 and 0 - p, where a: Nat and p: Pos
       preds \_\leq\_, \_\geq\_, \_<\_, \_>\_: Int \times Int;
               even, odd : Int
               -\_, sign : Int \rightarrow Int;
       ops
               abs : Int \rightarrow Nat;
               \_+\_, \_\times\_, \_-\_, min, max : Int \times Int \rightarrow Int;
               \_\hat{}: Int \times Nat \rightarrow Int;
               \_/?\_, \_div\_, \_quot\_, \_rem\_: Int \times Int \rightarrow? Int;
               \_mod\_: Int \times Int \rightarrow ? Nat
```

end

The rationals are specified as fractions of integers.

The specification RAT of rational numbers follows the same scheme as the specification of integers discussed above. This time, the specification INT is imported. The rationals are then defined as equivalence classes of pairs consisting of an integer and a positive number written as fractions, using the axiom:

 $\forall i, j: Int; p, q: Pos$ $\bullet i / p = j / q \Leftrightarrow i \times q = j \times p$ %(equality_Rat)%

The sort *Int* is then declared to be a subsort of *Rat*. Note that thanks to the behavior of subsorted overloading in CASL, the declaration of the operation:

 $_/_: Rat \times Rat \rightarrow ? Rat;$

allows rationals to be written also as x/y, for arbitrary integers x and $y \neq 0$.

spec RAT = INT then generated type $Rat ::= _/_(Int; Pos)$ sort Int < Ratpreds $_\leq_, _<_, _\geq_, _\geq_, _>_: Rat \times Rat$ ops $__, abs : Rat \rightarrow Rat;$ $_+_, ___, _\times_, min, max : Rat \times Rat \rightarrow Rat;$ $_-/_: Rat \times Rat \rightarrow ? Rat;$ $_, _: Rat \times Int \rightarrow Rat$

end

12.2 Library Basic/StructuredDatatypes

This library provides specifications of the familiar structured datatypes as used e.g. for the design of algorithms or within programming languages. Its main focus is data structures like (finite) sets, lists, strings, (finite) maps, (finite) bags, arrays, and various kinds of trees. Common to all these concepts is that they are generic. Consequently, all of the specifications of this library are generic.

Finite sets, maps and bags are specified as generated datatypes, with equality determined by means of observers.

Finite sets, finite maps and finite bags are specified using a generated sort. An observer operation or predicate is then introduced in order to define equality on this sort. Concerning finite sets, equality on the sort Set[Elem] is characterized using the predicate $_eps_$ (displayed as ϵ) in the specification GENERATESET. This leads to the extensionality axiom:

• $M = N \Leftrightarrow \forall x : Elem \bullet x \in M \Leftrightarrow x \in N$ %(equality_Set)%

library BASIC/STRUCTUREDDATATYPES

spec GENERATESET [**sort** Elem] = generated type $Set[Elem] ::= \{\} \mid _+_(Set[Elem]; Elem)\}$ pred __ ϵ __: $Elem \times Set[Elem]$ %% a system of representatives for sort Set[Elem] is %% %% {} and {} + x_1 + x_2 + ... + x_n %% %% where x_1 < x_2 < ... < x_n, n >= 1, x_i of type Elem, %% for an arbitrary strict total order _<_ on Elem. end spec Set [sort *Elem*] given NAT =GENERATESET [sort *Elem*] then %def **preds** *isNonEmpty* : *Set*[*Elem*]; $_\subseteq_: Set[Elem] \times Set[Elem]$ $\{_\}: Elem \rightarrow Set[Elem];$ ops \sharp .: Set[Elem] \rightarrow Nat; $_+_: Elem \times Set[Elem] \rightarrow Set[Elem];$ $___: Set[Elem] \times Elem \rightarrow Set[Elem];$ $__\bigcirc_, _\cup_, ____, _symDiff_:$ $Set[Elem] \times Set[Elem] \rightarrow Set[Elem]$

end

Finite maps, i.e. elements of the sort Map[S, T], are considered to be identical if looking up any value in S yields the same result in both cases:

```
• M = N \Leftrightarrow \forall s : S \bullet lookup(s, M) = lookup(s, N) %(equality_Map)%
```

On top of this, the specification MAP adds e.g. predicates for elementhood (ϵ) as well as for determining the profile of a map $(f :: x \to y \text{ means that } f \text{ is a map from } x \text{ to } y)$.

The specification TOTALMAP restricts maps to everywhere-defined maps (which are isomorphic to tuples). Since maps are finite, totality is only possible for maps over finite argument sorts. The latter is specified in the specification FINITE, using a partial surjection from the natural numbers. Since this specification is rather unusual, we make an exception and also show its axioms.

```
spec GENERATEMAP [sort S] [sort T] =

generated type Map[S,T] ::= empty \mid \_[\_/\_](Map[S,T]; T; S)

op lookup : S \times Map[S,T] \rightarrow ? T

end
```

```
spec MAP [sort S] [sort T] given NAT =
      GENERATEMAP [sort S] [sort T]
and Set [sort S]
and Set [sort T]
then %def
      free type Entry[S,T] ::= [\__/\_](target:T; source:S)
      preds isEmpty : Map[S,T];
              -\epsilon_{-}: Entry[S,T] \times Map[S,T];
              \_::\_ \rightarrow \_: Map[S,T] \times Set[S] \times Set[T]
              \_+\_, \_-\_: Map[S,T] \times Entry[S,T] \rightarrow Map[S,T];
      ops
              \_-\_: Map[S,T] \times S \to Map[S,T];
              \_\_\_: Map[S,T] \times T \to Map[S,T];
              dom: Map[S,T] \rightarrow Set[S];
              range : Map[S,T] \rightarrow Set[T];
              \_\cup\_: Map[S,T] \times Map[S,T] \rightarrow ? Map[S,T]
end
spec FINITE [sort Elem] =
              NAT
      ł
                      f: Nat \rightarrow ? Elem
        then op
              • \forall x: Elem \bullet \exists n: Nat \bullet f(n) = x
                                                                   %(f_surjective)%
               • \exists n: Nat \bullet \forall m: Nat \bullet def f(m) \Rightarrow m < n \qquad \% (f_bounded)\%
      reveal Elem
end
spec TOTALMAP [FINITE [sort S]] [sort T] =
               MAP [sort S] [sort T]
      {
        then sort
                       TotalMap[S,T] =
                           \{M: Map[S,T] \bullet \forall x: S \bullet def \ lookup(x, M)\}
                       [-/]: TotalMap[S,T] × T × S \rightarrow TotalMap[S,T];
               ops
                       lookup: S \times TotalMap[S,T] \rightarrow T;
                       \_+\_: TotalMap[S,T] \times Entry[S,T] \rightarrow TotalMap[S,T];
                       range : TotalMap[S,T] \rightarrow Set[T];
                       \_\cup\_: TotalMap[S,T] \times TotalMap[S,T]
                                                          \rightarrow? TotalMap[S,T]
              pred __\epsilon_ : Entry[S,T] × TotalMap[S,T]
      hide Map[S,T]
```

end

In the specification GENERATEBAG, those elements of sort Bag[Elem] are identified that show the same number of occurrences (observed by the operation *freq*) for all entries:

•
$$M = N \Leftrightarrow \forall x : Elem \bullet freq(M, x) = freq(N, x)$$
 %(equality_Bag)%

Lists are specified as a free datatype.

In the specification GENERATELIST, lists are built up from the empty list by adding elements in front. The usual list operations are provided: *first* and *last* select the first or last element of a list, while *rest* or *front* select the remaining list; $\#_{--}$ counts the number of elements in a list, while *freq* counts the number of occurrences of a given element; *take* takes the first *n* elements of a list, while *drop* drops them.

```
spec GENERATELIST [sort Elem] =
    free type List[Elem] ::= [] | _::_(first:?Elem; rest:?List[Elem])
end
```

end

Arrays are specified as certain finite maps.

The specification ARRAY includes the condition $\min \leq \max$ as an axiom in its first parameter. This ensures a non-empty index set. Arrays are defined as finite maps from the sort *Index* to the sort *Elem*, where the typical array operations lookup (__!_) and assignment (_!_ := __) are introduced in terms of finite map operations. Finally, revealing the essential signature elements yields the desired datatype.

```
spec ARRAY [ops min, max : Int • min \leq max \% (Cond_nonEmptyIndex)\%]
[sort Elem]
given INT =
sort Index = {i: Int • min <math>\leq i \land i \leq max}
then {
    MAP [sort Index] [sort Elem]
    with sort Map[Index,Elem] \mapsto Array[Elem],
    op empty \mapsto init
    then ops __!_:=_: : Array[Elem] \times Index \times Elem \rightarrow Array[Elem];
    __!_: : Array[Elem] \times Index \rightarrow ? Elem
}
reveal sort Array[Elem], ops init, _!_., _!_:=_
end
```

Several kinds of tree are available, differing in the branching and in the positions of elements.

The library concludes with several specifications concerning trees. There are specifications of binary trees (BINTREE, BINTREE2), k-branching trees (KTREE), and trees with possibly different branching at each node (NTREE). Each of these branching structures can be equipped with data in different ways: Either all nodes of a tree carry data (as is the case in BINTREE, KTREE, and NTREE), or just the leaves of a tree have a data entry (as in BINTREE2).

Binary trees admit two children for each internal node.

```
spec GENERATEBINTREE [sort Elem] =
     free type
          BinTree[Elem] ::= nil
                            | binTree(entry:?Elem; left:?BinTree[Elem];
                                     right:?BinTree[Elem])
end
spec BINTREE [sort Elem] given NAT =
     GENERATEBINTREE [sort Elem] and SET [sort Elem]
then preds isEmpty, isLeaf : BinTree[Elem];
            isCompoundTree : BinTree[Elem];
            \_\epsilon\_: Elem \times BinTree[Elem]
            height: BinTree[Elem] \rightarrow Nat;
     ops
            leaves: BinTree[Elem] \rightarrow Set[Elem]
end
spec GENERATEBINTREE2 [sort Elem] =
     free type NonEmptyBinTree2[Elem] ::=
                     leaf(entry:?Elem)
                    | binTree(left:?NonEmptyBinTree2[Elem];
                              right:?NonEmptyBinTree2[Elem])
     free type BinTree2[Elem] ::= nil | sort NonEmptyBinTree2[Elem]
end
spec BINTREE2 [sort Elem] given NAT =
     GENERATEBINTREE2 [sort Elem] and SET [sort Elem]
then %def
     preds isEmpty, isLeaf : BinTree2[Elem];
            isCompoundTree : BinTree2[Elem];
            \_\epsilon\_: Elem \times BinTree2[Elem]
            height : BinTree2[Elem] \rightarrow Nat;
     ops
            leaves : BinTree2[Elem] \rightarrow Set[Elem]
end
```

k-trees admit k children for each internal node (with k fixed).

We now come to k-branching trees. The branching is specified by using arrays of trees of size k, which are used to contain the children of a node in the tree.

```
spec GENERATEKTREE [op k : Int • k \ge 1 %(Cond_nonEmptyBranching)%]
                          [sort Elem] given INT =
      ARRAY [ops 1 : Int; k : Int
                   fit ops min : Int \mapsto 1, max : Int \mapsto k
               [sort KTree[k,Elem]]
then free type
           KTree[k, Elem] ::= nil
                              | kTree(entry:?Elem;
                                      branches:?Array[KTree[k,Elem]])
end
spec KTREE [op k : Int \bullet k \ge 1]
                                               %(Cond_nonEmptyBranching)%]
               [sort Elem]
      given INT =
      GENERATEKTREE [op k : Int] [sort Elem]
and SET [sort Elem]
then %def
      preds isEmpty, isLeaf : KTree[k,Elem];
             isCompoundTree : KTree[k, Elem];
             \_\epsilon\_: Elem \times KTree[k, Elem]
      ops
             height: KTree[k, Elem] \rightarrow Nat;
             maxHeight: Index \times Array[KTree[k, Elem]] \rightarrow Nat;
             leaves : KTree[k, Elem] \rightarrow Set[Elem];
             allLeaves : Index \times Array[KTree[k, Elem]] \rightarrow Set[Elem]
```

end

n-trees admit arbitrary branching.

Finally, *n*-trees are trees with possibly different branching at each node. This is specified by equipping each node in a tree with a list of child trees.

end

end

Case Study: The Steam-Boiler Control System

In this chapter we illustrate the use of CASL on a fairly large and complex case study, the steam-boiler control system. This case study is particularly interesting since it has been used several times as a competition problem, and many other specification frameworks have been illustrated with it, see [1]. Here we describe how to derive a CASL specification of the steam-boiler control system, starting from the informal requirements provided to the participants of the Dagstuhl meeting Methods for Semantics and Specification, organized jointly by Jean-Raymond Abrial, Egon Börger and Hans Langmaack in June 1995. The aim of this formalization process is to analyze the informal requirements, to detect inconsistencies and loose ends, and to translate the requirements into a CASL specification. During this process we have to provide interpretations for the unclear or missing parts. We explain how we can keep track of these additional interpretations by localizing very precisely in the formal specification where they lead to specific axioms, thereby taking care of the traceability issues. We also explain how the CASL specification is obtained in a stepwise way by successive analysis of various parts of the problem description. Finally we discuss the validation of the CASL requirements specification resulting from the formalization process, and in a last step we refine the requirements specification in a sequence of architectural specifications that describe the intended architecture of the steam-boiler control system.¹

The reader not already familiar with the steam-boiler control system case study may want to start by reading App. C, where the original description of the problem is reproduced.

¹ This chapter partially relies on an earlier work published in [8] where the PLUSS specification language [7, 9] was used together with the LARCH prover [25]. However, the specification methodology illustrated in this chapter is significantly improved, and moreover CASL provides several features that lead to a much more concise and perspicuous specification, as illustrated later in this chapter.

[©] IFIP International Federation for Information Processing 2004

13.1 Introduction

The aim of this chapter is to illustrate how one can solve the steam-boiler control specification problem with CASL. For this we have to provide a CASL specification of the software system that controls the level of water in the steam-boiler. Our work plan can be described as follows:

- 1. The main task is to derive, starting from the informal requirements, a requirements specification, written in CASL, of the steam-boiler control system. In particular, this task involves the following activities:
 - a) We must perform an in-depth analysis of the informal requirements. Obviously, this is necessary to gain a sufficient understanding of the problem to be specified, and this preliminary task may not seem worth mentioning. Let us stress, however, that the kind of preliminary analysis required for writing a formal specification proves especially useful to detect discrepancies in the informal requirements that would otherwise be very difficult to detect. Indeed, from our practical experience, this step is usually very fruitful from an engineering point of view, and one could argue that the benefits to be expected here are enough in themselves to justify the use of formal methods, even if for lack of time (or other resources) no full formal development of the system is performed.
 - b) Once we have a sufficient understanding of the problem to be specified, we must translate the informal requirements into a formal specification. This step will require us to provide interpretations for the unclear or missing parts of the informal requirements. Moreover, this formalization process will also be helpful to further detect inconsistencies and loose ends in the informal requirements. Here, a very important issue is to keep track of the interpretations made during the formalization process, in order to be able, later on, to take into account further modifications and changes of the informal requirements.
 - c) When we have written the formal requirements specification, we must carefully check its adequacy with respect to the informal requirements: this part is called the validation of the formal specification.

In principle there should be some interaction between the specification team and the team who has designed the informal requirements, in particular to check whether the suggested interpretations of the detected loose ends are adequate. In the framework of this case study, however, such interactions were not possible, and we can only use our intuition to assess the soundness of the interpretations made during the writing of the formal specification.

2. Once a validated requirements specification is obtained, we can proceed toward a program by a sequence of refinements. Here a crucial step is the choice of an architecture of the desired implementation, expressed by an architectural specification as explained in Chap. 8. Each refinement step leads to proof obligations which allow the correctness of the performed refinement to be assessed. In a last step, a program is derived from the final design specification.

Before starting to explain how to write a CASL requirements specification of the steam-boiler control system, let us make a few comments on this case study. First, note that, although in principle a hybrid system, the steam-boiler control system turns out to be merely a reactive system, not even a 'hard realtime' system (see e.g. the assumptions made in App. C.3). Moreover, even if the whole system, i.e., the control program and its physical environment is distributed, this is not the case, at least at the requirements level, for the steam-boiler control system. CASL turns out to be especially well-suited to capture the features of systems like the steam-boiler control system, where data and control are equally important (in particular, here data play a prominent role in failure detection). The various constructs provided by CASL allow the specifications to be formulated straightforwardly and perspicuously – and significantly more concisely than in other algebraic specification languages.

As a last remark we must make clear that for the sake of simplicity the *initialization* phase of the steam-boiler control system (see App. C.4.1) is not specified. However, it should be clear that it would be straightforward to extend our specification so as to take the initialization phase into account, following exactly the same methodology as for the rest of the case study.

This chapter is organized as follows. In Sect. 13.2 we start by providing some elementary specifications that will be useful for the rest of the case study. In Sect. 13.3 we explain how we will proceed to derive the CASL requirements specification in a stepwise way. Then in Sect. 13.4 we detail the specification of the mode of operation of the steam-boiler control system. In Sect. 13.5 we specify the detection of the various equipment failures, and in Sect. 13.6 we explain how we can compute, at each cycle, some predicted values for the messages to be received at the next cycle. In Sect. 13.9 we explain how our CASL requirements specification can be validated, and in Sect. 13.10 we refine the CASL requirements specification in a sequence of architectural specifications that describe the intended architecture of the implementation of the steam-boiler control system.

13.2 Getting Started

As explained in App. C.3, in each cycle the steam-boiler control system collects the messages received, performs some analysis of the information contained in them, and then sends messages to the physical units. We will therefore start with the specification of some elementary datatypes, such as "messages received" and "messages sent". To specify the messages sent and received, we follow App. C.5 and C.6. Note that some messages have parameters (e.g. pump number, pump state, pump controller state, mode of operation), and we must therefore specify the corresponding datatypes as well. For the sake of clarity, we group together all similar messages (e.g. all "repaired" messages, all "failure acknowledgement" messages) by introducing a suitable parameter "physical unit". A physical unit is either a pump, a pump controller, the water level measuring device or the steam output measuring device. Remember that we do not specify the physical units as such, since we do not specify the physical environment of the steam-boiler (we do not specify the steam-boiler either, we only specify the steam-boiler control system). Hence the datatype "physical unit" is just an elementary datatype that says that we have some pumps, some pump controllers, and the two measuring devices.

Some messages have a value v as parameter. From the informal requirements we can infer that these values are (approximations of) real numbers, but it is not necessary at this level to make any decision about the exact specification of these values. In our case study, we will therefore rely on a very abstract (loose) specification VALUE, introducing a sort *Value* together with some operations and predicates, which are left unspecified (we expect of course that these operations and predicates will have the intuitive interpretation suggested by their names). This means that we consider VALUE as being a general parameter of our specification.² This point is discussed again in Sect. 13.10. Note also that we will abstract from measuring units (such as liter, liter/sec), since ensuring that these units are consistently used is a very minor aspect of this particular case study.³

This first analysis leads to the following specifications: VALUE, BASICS, MESSAGES_SENT, and MESSAGES_RECEIVED.

from BASIC/NUMBERS get NAT

```
%display _half %LATEX _/2
%display _square %LATEX _<sup>2</sup>
```

then sorts Nat < Valueops $...+.: Value \times Value \rightarrow Value, assoc, comm, unit 0;$ $...-.: Value \times Value \rightarrow Value;$ $...\times.: Value \times Value \rightarrow Value, assoc, comm, unit 1;$

 $^{^2}$ We leave VALUE as an implicit parameter of our specifications, rather than using generic specifications taking VALUE as a parameter, since our specifications are not to be instantiated by argument specifications describing several kinds of values, but on the contrary should all refer to the same abstract datatype of values.

³ It is of course possible to take measuring units into account, following for instance the method described in [18]. Appropriate CASL libraries supporting measuring units are currently being developed.

```
\_/2, \_^2: Value \rightarrow Value;
            min, max : Value \times Value \rightarrow Value
     preds \_ < \_, \_ \le \_: Value × Value
end
spec BASICS =
     free type PumpNumber ::= Pump1 | Pump2 | Pump3 | Pump4;
     free type PumpState ::= Open | Closed;
     free type PumpControllerState ::= Flow | NoFlow;
     free type PhysicalUnit ::= Pump(PumpNumber)
                              PumpController(PumpNumber)
                              SteamOutput | WaterLevel;
     free type Mode ::= Initialization | Normal | Degraded
                       Rescue | EmergencyStop;
end
spec Messages_Sent =
     BASICS
then free type
     S\_Message ::= MODE(Mode) | PROGRAM\_READY | VALVE
                   OPEN_PUMP(PumpNumber)
                   CLOSE_PUMP(PumpNumber)
                  FAILURE_DETECTION(PhysicalUnit)
                  REPAIRED_ACKNOWLEDGEMENT(PhysicalUnit);
end
spec Messages_Received = 
     BASICS and VALUE
then free type
     R\_Message ::= STOP \mid STEAM\_BOILER\_WAITING
                   PHYSICAL_UNITS_READY
                   PUMP_STATE(PumpNumber; PumpState)
                   PUMP_CONTROLLER_STATE(PumpNumber;
                                                PumpControllerState)
                   LEVEL(Value) | STEAM(Value)
                   REPAIRED(PhysicalUnit)
                   FAILURE_ACKNOWLEDGEMENT(PhysicalUnit)
                   junk;
```

end

In addition to the "messages received" specified in App. C.6, we add an extra constant message *junk*. This message will represent any message received which does not belong to the class of recognized messages. We do not add a similar message to the messages sent, since we may assume that the steamboiler control system will only send proper messages. Obviously, receiving a

junk message will lead to the detection of a failure of the message transmission system.

In the SBCS_CONSTANTS specification we describe the various constants that characterize the steam-boiler (these constants are explained in App. C.2.6).

```
spec SBCS_CONSTANTS =
VALUE
then ops C, M1, M2, N1, N2, W, U1, U2, P : Value;
dt : Value \%\% Time duration between two cycles (5 sec.)
\%\% These constants must verify some obvious properties:
• 0 < M1 • M1 < N1 • N1 < N2 • N2 < M2 • M2 < C
• 0 < W • 0 < U1 • 0 < U2 • 0 < P
```

end

We will also specify the datatypes "set of messages received" and "set of messages sent" since, as suggested at the end of App. C.3, all messages are supposed to be received (or emitted) simultaneously at each cycle. The two latter specifications are obtained by instantiation of a generic specification SET of "sets of elements", which is imported from the library BA-SIC/STRUCTUREDDATATYPES.

from BASIC/STRUCTUREDDATATYPES get SET

```
spec PRELIMINARY =
SET [MESSAGES_RECEIVED fit Elem \mapsto R\_Message]
and SET [MESSAGES_SENT fit Elem \mapsto S\_Message]
and SBCS_CONSTANTS
end
```

 \heartsuit As illustrated by the above specifications, it is particularly convenient to *structure* our formal specification into coherent, easy to grasp, *named specifications* that will be easily reused later on by referring to their names (and this of course will prove even more important in the sequel). Moreover, *free datatypes* are especially useful here to obtain concise specifications. On the other hand, loose specifications are useful to avoid overspecification of values in VALUE and of the steamboiler constants in SBCS_CONSTANTS. Declaring that Nat is a subsort of Value ensures that natural numbers can be used as arguments of operations on values. Reusing standard specifications of usual datatypes from the BASIC *libraries* avoids the need to specify them again, and of course it is essential that these specifications are *generic* in order to easily adapt them as desired when they are reused. Finally, *display* annotations are useful to conveniently display some symbols as usual mathematical symbols.⁴ \heartsuit

⁴ In this chapter, metacomments about the adequacy of CASL features will be highlighted like this.

13.3 Carrying On

As emphasized in App. C.3, the steam-boiler control system is a typical example of a control-command system. The specification of such systems always follows the same pattern:

- A preliminary set of specifications group all the basic information about the system to be controlled, such as the specification of the various messages to be exchanged between the system and its environment, and the specification of the various constants related to the system of interest. This is indeed the aim of the specification PRELIMINARY introduced in the previous section.
- Then, the various states of the control system should be described. At this stage, however, it would be much too early to determine which state variables are needed. Thus states will be represented by values of a (loosely specified) sort *State*, equipped with some observers (corresponding to access to state variables). During the requirements analysis and formalization phase we may need further observers, to be introduced on a by-need basis.
- Then a (group of) specification(s) will take care of the analysis of the messages received here, of failure detection in particular. On the basis of this analysis, some actions should be taken, corresponding to the messages to be sent to the environment. State variables are also updated according to the result of the analysis of the messages received and to the messages to be sent.
- Finally a specification describes the overall control-command system as a labeled transition system.

A very rough preliminary sketch of the steam-boiler control system specification looks therefore as follows:

```
library SBCS

from BASIC/NUMBERS get NAT

from BASIC/STRUCTUREDDATATYPES get SET

%display _half %LATEX _/2

%display _square %LATEX _2

:

spec PRELIMINARY = %{ See previous section. }%

spec SBCS_STATE =

PRELIMINARY

then sort State

ops %% Needed state observers are introduced here.

%% E.g., we need an observer for the mode of operation:

mode : State \rightarrow Mode;

...
```

end

```
spec SBCS_ANALYSIS =
       SBCS_STATE
then %% Analysis of messages received and in particular failure detection.
       %% Computation of the messages to be sent.
       op messages_to_send : State \times Set[R\_Message] \rightarrow Set[S\_Message];
       %% Computation of the updates of the state variables.
       % For each observer obs defined in SBCS_STATE,
           we introduce an operation next_obs that computes the
           corresponding update according to the analysis of the messages
           received in this round. For instance, we specify here an operation
           next_mode corresponding to the update of the observer mode. \%
       ops next_mode : State \times Set[R_Message] \rightarrow Mode;
end
       STEAM_BOILER_CONTROL_SYSTEM =
spec
       SBCS_ANALYSIS
then op
              init : State
       pred is\_step: State \times Set[R\_Message] \times Set[S\_Message] \times State
       \%\% Specification of the initial state init by means of the observers, e.g.
        • mode(init) = \dots
        • . . .
       % Specification of is_step by means of the observers
           and of the updating operations, e.g.: \}\%
       \forall s, s' : State; msgs : Set[R_Message]; Smsg : Set[S_Message]
        • is\_step(s, msqs, Smsq, s') \Leftrightarrow
             mode(s') = next\_mode(s, msgs) \land \ldots \land
             Smsq = messages\_to\_send(s, msqs)
then %% Specification of the reachable states:
       free { pred reach : State
               \forall s, s' : State; msgs : Set[R_Message]; Smsg : Set[S_Message]
                • reach(init)
                • reach(s) \land is\_step(s, msgs, Smsg, s') \Rightarrow reach(s')
```

end

Of course the specification SBCS_ANALYSIS is likely to be further structured into smaller pieces of specifications. Indeed, since the informal requirements are too complex to be handled as a whole, we will therefore successively concentrate on various parts of them. The study and formalization of each chunk of requirements will lead to specifications that will later on be put together to obtain the SBCS_ANALYSIS specification. As already pointed out, it is likely that when analyzing a chunk of requirements we will discover the need for new observers on states (i.e., new state variables). This means that the specification SBCS_STATE will be subject to iterated extensions where we introduce the new observers that are needed.

}

For instance, in App. C.6 it is explained that when the STOP message has been received three times in a row, the program must go into the *EmergencyStop* mode. We need therefore an observer (i.e., a state variable) to record the number of times we have successively received the STOP message. So in the sequel we will start from the following specification of states:

```
spec SBCS_STATE_1 =
PRELIMINARY
then sort State
ops mode : State \rightarrow Mode;
numSTOP : State \rightarrow Nat
```

end

Introducing the new observer *numSTOP* means that we will have to specify a corresponding *next_numSTOP* operation in the SBCS_ANALYSIS specification.

 \heartsuit Let us insist again on the importance of *structuring* our formal specification into coherent, easy to grasp, *named specifications* that will be easy to reuse later on by referring to their names. As explained above it is moreover essential to rely on a *loose specification* of states so that we can introduce later on as many observers as needed. Using a *predicate* (such as *is_step*) to describe a labeled transition system is quite convenient here, and provides us with an elegant way of handling both input and output for each transition. Finally, it is essential to use a *free constraint* to specify the reachable states, and thus we need to *combine parts with a loose interpretation and parts with an initial interpretation in the same specification.*

13.4 Specifying the Mode of Operation

Our next step is the specification of the various operating modes in which the steam-boiler control system operates. (As explained in Sect. 13.1 we do not take into account the *Initialization* mode in this specification.) According to App. C.4, the operating mode of the steam-boiler control system depends on which failures have been detected (see e.g. "all physical units [are] operating correctly", "a failure of the water level measuring unit", "detection of an erroneous transmission"). It depends also on the expected evolution of the water level (see "If the water level is risking to reach…").

We will therefore assume that the specification SBCS_ANALYSIS will provide the following predicates which, given a known state and newly received messages, should reflect the failures detected by the steam-boiler control system: 5

⁵ It is important to make a subtle distinction between the actual failures, about which we basically know nothing, and the failures detected by the steam-boiler

- Transmission_OK : State × Set[R_Message] should hold iff we rely on the message transmission system,
- $PU_OK : State \times Set[R_Message] \times PhysicalUnit$ should hold iff we rely on the corresponding physical unit,
- Dangerous WaterLevel : $State \times Set[R_Message]$ should hold iff we estimate that the water level risks reaching the min (M1) or max (M2) limits.

However, at this stage our understanding of the steam-boiler control system is still quite preliminary, and it is therefore too early to attempt to specify these predicates. Therefore, our specification MODE_EVOLUTION, where we specify the new operating mode according to the previous one and the newly received messages (i.e., the operation *next_mode*), will be made *generic* w.r.t. these predicates. Let us emphasize that here genericity is used to ensure a *loose coupling* between the current specification of interest, MODE_EVOLUTION, and other specifications expected to provide the needed predicates.

Let us now explain how to specify the new mode of operation. At first glance the informal requirements (see App. C.4) look quite complicated, mainly because they explain, for each operating mode, under which conditions the steam-boiler control system should stay in the same operating mode or switch to another one. However, things get simpler if we analyze under which conditions the next mode is one of the specified operating modes. In particular, a careful analysis of the requirements shows that, except for switching to the *EmergencyStop* mode, we can determine the new operating mode (after receiving some messages) without even taking into account the previous one.

To improve the legibility of our specification it is better to introduce some auxiliary predicates (*Everything_OK*, *AskedToStop*, *SystemStillControllable*, and *Emergency*) that will facilitate the characterization of the conditions under which the system switches from one mode to another:

- The aim of the predicate $Everything_OK$ is to express that we believe that all physical units are operating correctly, including the message transmission system.
- The aim of the predicate *AskedToStop* is to determine if we have received the *STOP* message three times in a row.
- The aim of the predicate *SystemStillControllable* is to characterize the conditions under which the steam-boiler control system will operate in *Rescue* mode. Let us point out that the corresponding part of the informal requirements (see App. C.4.4) is not totally clear, in particular the exact meaning of the sentence "if one can rely upon the information which comes from the units for controlling the pumps". There is a double ambiguity here: on the one hand it is unclear whether "the pumps" means "all pumps" or "at least

control system. The behavior of the steam-boiler control system is induced by the failures detected, whatever the actual failures are.

one pump"; on the other hand there are two ways of "controlling" each pump (the information sent by the pump and the information sent by the pump controller), and it is unclear whether "controlling" refers to both of them or only to the pump controller. Our interpretation will be as follows: we consider it is enough that at least one pump is "correctly working", and for us correctly working will mean we rely on both the pump and the associated pump controller. As with all interpretations made during the formalization process, we should in principle interact with the designers of the informal requirements in order to clarify what was the exact intended meaning and to check that our interpretation is adequate. The important point is that our interpretation is entirely localized in the axiomatization of *SystemStillControllable*, and it will therefore be fairly easy to change our specification in case of misinterpretation.

• The aim of the predicate *Emergency* is to characterize when we should switch to the *EmergencyStop* mode. In App. C.4.2, it is said that the steam-boiler control system should switch from *Normal* mode to *Rescue* mode as soon as a failure of the water level measuring unit is detected. However, in App. C.4.4, it is explained that the steam-boiler control system can only operate in *Rescue* mode if some additional conditions hold (represented by our predicate *SystemStillControllable*). We decide therefore that when in *Normal* mode, if a failure of the water level measuring unit is detected, the steam-boiler control system will switch to *Rescue* mode only if *SystemStillControllable* holds, otherwise it will switch (directly) to *EmergencyStop* mode.⁶

The axiomatization of the next mode of operation is now both simple and clear, as illustrated by the MODE_EVOLUTION specification.⁷

local %% Auxiliary predicates to structure the specification of *next_mode*.

⁶ If our interpretation is incorrect, then in some cases we may have replaced a sequence Normal \rightarrow Rescue \rightarrow EmergencyStop by a sequence Normal \rightarrow EmergencyStop. Note that a sequence Normal \rightarrow Rescue \rightarrow Normal or Degraded is not possible since several cycles are necessary between a failure detection and the decision that the corresponding unit is again fully operational, see Sect. 13.5, i.e., we must have a sequence of the form Normal \rightarrow Rescue $\rightarrow \ldots \rightarrow$ Rescue \rightarrow Normal or Degraded in such cases.

⁷ Note that once in the *EmergencyStop* mode, we specify that we stay in this mode forever, rather than specifying that the steam-boiler control system actually stops. Note also that we realize that the operation *next_numSTOP* is better specified in this MODE_EVOLUTION specification.

```
preds Everything_OK, AskedToStop, SystemStillControllable,
                Emergency: State \times Set[R\_Message]
       \forall s: State; msqs: Set[R\_Message]
        • Everything_OK(s, msqs) \Leftrightarrow
              (Transmission_OK(s, msqs) \land
                    (\forall pu : PhysicalUnit \bullet PU\_OK(s, msgs, pu)))
        • AskedToStop(s, msgs) \Leftrightarrow numSTOP(s) = 2 \land STOP \ \epsilon \ msgs
        • SystemStillControllable(s, msgs) \Leftrightarrow
              (PU\_OK(s, msgs, SteamOutput) \land
             (\exists pn : PumpNumber \bullet PU\_OK(s, msqs, Pump(pn)))
                                     \land PU\_OK(s, msgs, PumpController(pn))))
        • Emergency(s, msgs) \Leftrightarrow
             (mode(s) = EmergencyStop \lor
               AskedToStop(s, msqs) \lor
               \neg Transmission_OK(s, msgs) \lor
               DangerousWaterLevel(s, msqs) \lor
               (\neg PU\_OK(s, msqs, WaterLevel) \land
                 \neg SystemStillControllable(s, msqs)))
within ops next\_mode: State \times Set[R\_Message] \rightarrow Mode;
              next\_numSTOP: State \times Set[R\_Message] \rightarrow Nat
       \forall s : State; msgs : Set[R\_Message]
       %% Emergency stop mode:
        • Emergency(s, msgs) \Rightarrow next\_mode(s, msgs) = EmergencyStop
       %% Normal mode:
        • \neg Emergency(s, msqs) \land
           Everything_OK(s, msgs) \Rightarrow next_mode(s, msgs) = Normal
       %% Degraded mode:
        • \neg Emergency(s, msqs) \land
          \neg Everything_OK(s, msqs) \land
           PU_OK(s, msgs, WaterLevel) \land
           Transmission_OK(s, msgs) \Rightarrow next_mode(s, msgs) = Degraded
       %% Rescue mode:
        • \neg Emergency(s, msqs) \land
           \neg PU\_OK(s, msqs, WaterLevel) \land
           SystemStillControllable(s, msqs) \land
           Transmission_OK(s, msgs) \Rightarrow next_mode(s, msgs) = Rescue
       %% next_numSTOP:
        • next\_numSTOP(s, msqs) = numSTOP(s) + 1 when STOP \epsilon msqs
                                          else 0
```

\mathbf{end}

In the next step of our formalization process, we will specify the predicates assumed by MODE_EVOLUTION, which amounts to specifying the detection of equipment failures. This will be the topic of the next section.

 \heartsuit Two essential features of CASL have been used in the specification MODE_EVOLUTION. On the one hand, the use of a *generic specification* (with imports) ensures loose coupling of the current specification of interest with the rest of the steam-boiler control system specification. On the other hand, *auxiliary predicates* improve the legibility of the specification, and declaring them in the *local part* of the specification ensures they are hidden and therefore not exported. \heartsuit

13.5 Specifying the Detection of Equipment Failures

The detection of equipment failures is described in App. C.7. It is quite clear that this detection is the most difficult part to formalize, mainly because both our intuition and the requirements (see e.g. "knows from elsewhere", "incompatible with the dynamics") suggest that we should take into account some inter-dependencies when detecting the various possible failures.

For instance, if we ask a pump to stop, and if in the next cycle the pump state still indicates that the pump is open, we may in principle infer either a failure of the message transmission system (e.g. the stop order was not properly sent or was not received, or the message indicating the pump state has been corrupted) or a failure of the pump (which was not able to execute the stop order or which sends incorrect state messages). Our understanding of the requirements is that in such a case we must conclude there has been a failure of the pump, not of the message transmission system. Let us stress again that it is important to distinguish between the actual failures of the various pieces of equipment, and the diagnosis we will make. Only the latter is relevant in our specification.

13.5.1 Understanding the Detection of Equipment Failures

Before starting to specify the detection of equipment failures, we must proceed to a careful analysis of App. C.7, in order to clarify the inter-dependencies mentioned above. Only then will we be able to understand how to structure our specification of this crucial part of the problem.

A first rough analysis of the part of App. C.7 devoted to the description of potential failures of the physical units (i.e. of the pumps, the pump controllers and the two measuring devices) shows that these failures are detected on the basis of the information contained in the messages received: we must check that the received values are in accordance with some *expected* values (according to the history of the system, i.e. according to the "dynamics of the system" and to the messages previously sent by the steam-boiler control system). In particular, the detection of failures of the physical units relies on the fact that we have effectively received the necessary messages. If we have not received these messages, then we should conclude there has been a failure of the message transmission system (see below), and in these cases (see the MODE_EVOLUTION specification), the steam-boiler control system switches to the *EmergencyStop* mode. The further detection of failures of the physical units (in addition to the already detected failure of the message transmission system) is therefore irrelevant in such cases.

Let us now consider the message transmission system. The description of potential failures of the message transmission system in App. C.7 is quite short. Basically, it tells us that we should check that the steam-boiler control system has received all the messages it was expecting, and that none of the messages received is aberrant. However, it is important to note that the involved analysis of the messages received combines two aspects: on the one hand, there is some 'static' analysis of the messages received in order to check that all messages that must be present in each transmission are effectively present (see App. C.6). These messages are exactly the messages required to proceed to the detection of the failures of the physical units. On the other hand, the steam-boiler control system expects to receive (or, on the contrary, not to receive) some specific messages according to the history of the system (for instance, the steam-boiler control system expects to receive a "failure acknowledgement" from a physical unit once it has detected a corresponding failure and sent a "failure" message to this unit, but not before), and here some 'dynamic' analysis is required. Obviously, the static analysis of the messages can be made on the basis of the messages received only, while the dynamic analysis must take into account, in addition to the messages received, the history of the system, and more precisely the history of the failures detected so far and of the "failure acknowledgement" and "repaired" messages received so far.

From this first analysis we draw the following conclusions on how to specify the detection of equipment failures:

- 1. In a first step we should keep track of the failure status of the physical units. This will lead to a new observer *status* on states, and to a specification STATUS_EVOLUTION of how this status evolves, i.e., of a *next_status* operation.
- 2. Then we specify the detection of the message transmission system failures (hence $Transmission_OK$) in the specification MESSAGE_TRANSMISSION_SYSTEM_FAILURE. As explained above, in a first step we take care of the static analysis of the messages received, and then in a second step we take care of the dynamic analysis of the messages received, using how we have kept track of the "status" of the physical units, i.e., using the observer status.
- 3. Then, for each physical unit, we specify the detection of its failures by comparing the message received with the *expected* one. For this comparison we can freely assume that the static analysis of the messages received has been successful, i.e., that the message sent by the physical unit has been received.

The corresponding specifications are described in the next subsections.

13.5.2 Keeping Track of the Status of the Physical Units

Remember that to perform the dynamic analysis of the messages received, as explained above, we must check that we receive "failure acknowledgement" and "repaired" messages when appropriate. In order to do this, we must keep track of the failures detected and of the "failure acknowledgement" and "repaired" messages received. Since the same reasoning applies for all physical units, we can do the analysis in a generic way. For each physical unit, we will keep track of its status, which can be either OK, *Failure WithoutAck* or *Failure WithAck*. The status of a physical unit will then be updated accordingly to the detection of failures, and receipt of "failure acknowledgement" and "repaired" messages.

Thus, in a first step we should extend the specification SBCS_STATE_1 to add an observer related to the failure status of physical units:

```
spec SBCS_STATE_2 =
SBCS_STATE_1
then free type Status ::= OK \mid Failure WithoutAck \mid Failure WithAck
op status : State × PhysicalUnit \rightarrow Status;
```

end

Now the specification of how the status of a physical unit evolves, i.e., of the operation *next_status* in STATUS_EVOLUTION, is quite straightforward. We rely again on the predicate PU_OK .⁸

```
spec STATUS_EVOLUTION

[pred PU_OK : State \times Set[R_Message] \times PhysicalUnit]

given SBCS_STATE_2 =

op next\_status : State \times Set[R_Message] \times PhysicalUnit \rightarrow Status

\forall s : State; msgs : Set[R_Message]; pu : PhysicalUnit

• status(s, pu) = OK \land PU_OK(s, msgs, pu)

\Rightarrow next\_status(s, msgs, pu) = OK

• status(s, pu) = OK \land \neg PU_OK(s, msgs, pu)

\Rightarrow next\_status(s, msgs, pu) = Failure WithoutAck

• status(s, pu) = Failure WithoutAck \land

FAILURE\_ACKNOWLEDGEMENT(pu) \epsilon msgs

\Rightarrow next\_status(s, msgs, pu) = Failure WithAck

• status(s, pu) = Failure WithoutAck \land

\neg (FAILURE\_ACKNOWLEDGEMENT(pu) \epsilon msgs)
```

```
\Rightarrow next_status(s, msgs, pu) = Failure WithoutAck
```

⁸ The reader may detect that the specification STATUS_EVOLUTION is not completely correct. However, we prefer to give here the text of the specification as it was originally written, and we will explain in Sect. 13.9 how we detect, when validating the specification of the steam-boiler control system, that something is not correct, and how the problem can be fixed.

- $status(s, pu) = FailureWithAck \land REPAIRED(pu) \epsilon msgs$ $\Rightarrow next_status(s, msgs, pu) = OK$
- $status(s, pu) = FailureWithAck \land \neg (REPAIRED(pu) \ \epsilon \ msgs)$ $\Rightarrow next_status(s, msgs, pu) = FailureWithAck$

end

 \heartsuit Here again we rely on a *generic specification with imports* to ensure loose coupling. As claimed earlier, the *loose specification* of states makes it easy to introduce further observers (hence further state variables). \heartsuit

13.5.3 Detection of the Message Transmission System Failures

As explained above, we first specify the static analysis of the messages received, and then we specify the dynamic analysis of these messages.

To specify the static analysis of messages, it is necessary to check that all "indispensable" messages are present. In addition, a set of messages received is "acceptable" if there are no "duplicated" messages in this set. Since we have specified the collection of messages received as a set, we cannot have several occurrences of exactly the same message in this set. (Note that this means that our choice of using "sets" instead of "bags", for instance, is therefore not totally innocent: either we assume that receiving several occurrences of exactly the same message will never happen, and this is an assumption about the environment, or we assume that this case should not lead to the detection of a failure of the message transmission system, and this is an assumption about the requirements.) However, specifying the collection of messages received as a set does not imply that a set of messages received cannot contain several LEVEL(v) messages, with distinct values (for instance). Hence we must check this explicitly.

Remember that receiving "unknown" messages (i.e., messages that do not belong to the list of messages as specified in App. C.6) is taken into account via the extra constant *junk* message (see the specification MESSAGES_RECEIVED). Another erroneous situation is when we simultaneously receive a failure acknowledgement and a repaired message for the same physical unit, i.e., that at least one cycle is needed between acknowledging the failure and repairing the unit. We will check this as well.⁹

We focus now on the dynamic analysis of the messages received. As explained above, to perform this dynamic analysis, we check that we receive "failure acknowledgement" and "repaired" messages when appropriate, according to the current status of each physical unit. We understand that for

⁹ We must confess that this belief is induced by our intuition about the behavior of the system. Indeed nothing in the requirements allows us to make either this interpretation or the opposite one. Although not essential, this assumption will simplify the axiomatization.

each failure signaled by the steam-boiler control system, the corresponding physical unit will send just one failure acknowledgement. Moreover, we will specify the steam-boiler control system in such a way that when it receives a "repaired" message, the steam-boiler control system acknowledges it immediately. Hence, if there is no problem with the message transmission system, and due to the fact that transmission time can be neglected, the steam-boiler control system must in principle receive only one repaired message for a given failure. Note that this does not contradict the "until..." part of the sentences describing the "repaired" messages in the informal requirements (see App. C.6). To summarize, we consider that we have received an unexpected message when:

- the program receives initialization messages but is no longer in initialization mode; or
- the program receives for some physical unit a "failure acknowledgement" without having previously sent the corresponding failure detection message, or receives redundant failure acknowledgements; or
- the program receives for some physical unit a "repaired message", but the unit is OK or its failure is not yet acknowledged.

We now have all the ingredients required to specify the $Transmission_OK$ predicate, taking into account both static and dynamic aspects, which leads to the following MESSAGE_TRANSMISSION_SYSTEM_FAILURE specification.

```
Message_Transmission_System_Failure =
spec
        SBCS_STATE_2
then local %% Static analysis:
        pred __is_static_OK : Set[R\_Message]
        \forall msgs : Set[R\_Message]
        • msgs is_static_OK \Leftrightarrow
              (\neg(junk \ \epsilon \ msgs) \land
                (\exists !v : Value \bullet LEVEL(v) \epsilon msgs) \land
                (\exists !v : Value \bullet STEAM(v) \epsilon msgs) \land
                (\forall pn : PumpNumber \bullet \exists !ps : PumpState \bullet
                        PUMP\_STATE(pn, ps) \in msgs) \land
                (\forall pn : PumpNumber \bullet \exists ! pcs : PumpControllerState \bullet
                        PUMP\_CONTROLLER\_STATE(pn, pcs) \in msqs) \land
                (\forall pu : PhysicalUnit \bullet
                        \neg (FAILURE_ACKNOWLEDGEMENT(pu) \epsilon msqs
                           \wedge REPAIRED(pu) \in msqs)))
               %% Dynamic analysis:
       pred __is_NOT_dynamic_OK_for_ : Set[R_Message] \times State
```

Transmission_OK(s, msgs) ⇔
 (msgs is_static_OK ∧ ¬(msgs is_NOT_dynamic_OK_for s))

end

 \heartsuit Here again *auxiliary predicates* declared in the *local part* of the specification are quite useful to improve the legibility of the specification. Note also the use of nested *quantifiers* in axioms (' \forall ', ' \exists ' as well as ' \exists !') – without them the axioms would be much more intricate, or further auxiliary operations would be needed. \heartsuit

13.5.4 Detection of the Pump and Pump Controller Failures

We start by considering the detection of the failures of the pumps.

As explained in Sec 13.5.1, we rely on the predicted pump state message. Thus, in a first step we should extend the specification SBCS_STATE_2 to add an observer related to the prediction of pump state messages. The prediction (*Open* or *Closed*) can however only be made when the status of the corresponding pump is OK. This is why we extend the sort *PumpState* to introduce a constant *Unknown_PS*:

```
spec SBCS_STATE_3 =
	SBCS_STATE_2
then free type ExtendedPumpState ::= sort PumpState | Unknown_PS
	op PS_predicted : State × PumpNumber \rightarrow ExtendedPumpState;
	%{ status(s, Pump(pn)) = OK \Leftrightarrow
	\neg (PS_predicted(s, pn) = Unknown_PS)}%
```

end

The specification of the detection of pump failures is now straightforward and is given in the PUMP_FAILURE specification. Remember that the meaning of $Pump_OK$ is only relevant when $Transmission_OK$ holds, which in particular implies that for each pump, there is only one $PUMP_STATE$ message for it in *msgs*. Moreover, we check the received value only if the predicted value is not $Unknown_PS$.

end

Let us now consider the detection of the failures of the pump controllers. Again we rely on the predicted pump state controller message. Here, we must be a bit careful in order to reflect the fact that stopping a pump has an instantaneous effect, while starting it takes five seconds (see App. C.2.3). Since five seconds is, unfortunately, exactly the elapsed time between two cycles, when we decide to activate a pump we may have to wait two cycles to receive a corresponding *Flow* pump controller state. This is why, in addition to the constant *Unknown_PCS*, used for the cases where no prediction can be made since the pump controller is not working correctly, we also introduce a constant *SoonFlow* to be used for the prediction related to a just activated pump.

```
spec SBCS_STATE_4 =

SBCS_STATE_3

then free type

ExtendedPumpControllerState ::= sort PumpControllerState

| SoonFlow | Unknown_PCS

op PCS\_predicted : State \times PumpNumber

\rightarrow ExtendedPumpControllerState;

\%{ status(s, PumpController(pn)) = OK \Rightarrow

\neg (PCS\_predicted(s, pn) = Unknown\_PCS)}%
```

end

The specification of the detection of pump controller failures is now straightforward and is given in the PUMP_CONTROLLER_FAILURE specification. Remember that the meaning of $Pump_Controller_OK$ is only relevant when $Transmission_OK$ holds, which in particular implies that for each pump, there is only one $PUMP_CONTROLLER_STATE$ message for it in *msgs*. Moreover, we check the received value only if the predicted value is either *Flow* or *NoFlow*, since if it is *SoonFlow* or *Unknown_PCS* we cannot conclude.

```
spec PUMP_CONTROLLER_FAILURE = SBCS_STATE_4
```

then pred $Pump_Controller_OK : State \times Set[R_Message] \times PumpNumber$ $\forall s : State; msgs : Set[R_Message]; pn : PumpNumber$

```
    Pump_Controller_OK(s, msgs, pn) ⇔

        PCS_predicted(s, pn) = Unknown_PCS

        ∨ PCS_predicted(s, pn) = SoonFlow

        ∨ PUMP_CONTROLLER_STATE(pn,

        PCS_predicted(s, pn) as PumpControllerState) ∈ msgs
```

\mathbf{end}

 \heartsuit In the above specifications, using *supersorts* to extend previously defined datatypes is particularly convenient, and avoids the need to explicitly relate values of *PumpState* and values of *ExtendedPumpState* (and similarly for *PumpControllerState*). Note the use of explicit *castings* in the axioms – in particular, the fact that *predicates do not hold* on undefined arguments resulting from castings is used in the above specifications. \heartsuit

13.5.5 Detection of the Steam and Water Level Measurement Device Failures

To specify the failures of the steam and water level measurement devices, we must again rely on some predicted values. Here however we cannot predict an exact value, but only an interval in which the received value should be contained. This leads to the following extension of SBCS_STATE_4:

\mathbf{spec}	$SBCS_STATE_5 =$
	SBCS_STATE_4
\mathbf{then}	free type $Valpair ::= pair(low : Value; high : Value)$
	ops steam_predicted, level_predicted : State \rightarrow Valpair;
	$% \{ low(steam_predicted(s)) \text{ is the minimal steam output predicted}, \}$
	$high(steam_predicted(s))$ is the maximal steam output predicted,
	and similarly for $level_predicted$. $\}\%$

end

The specification of the failures of the measurement devices is again straightforward and is given in the STEAM_FAILURE and LEVEL_FAILURE specifications. Remember that the meaning of $Steam_OK$ (Level_OK resp.) is only relevant when Transmission_OK holds, which in particular implies that there is only one STEAM(v) (LEVEL(v) resp.) message in msgs (hence only one possible v in the quantifications $\forall v : Value \ldots$ below). Note also that here we assume that the predicted values will take care of the static limits (θ and W for the steam, θ and C for the water level), thus we do not need to check these static limits explicitly here.

end

13.5.6 Summing Up

We now have all the ingredients necessary for the specification of the predicate PU_OK . This is done in the FAILURE_DETECTION specification, which integrates together all the specifications related to failure detection.

 $(low(level_predicted(s)) \le v) \land$ $(v \le high(level_predicted(s)))))$

} hide ops Pump_OK, Pump_Controller_OK, Steam_OK, Level_OK
end

 \heartsuit In the above specification, we rely on explicit *hiding* of operations that are no longer needed. Moreover, the *'same name, same thing' principle* is essential here: each of the five specifications extended in FAILURE_DETECTION is itself an extension of some specification SBCS_STATE_*i* of states, but with the 'same name, same thing' principle we get the effect that each of them extends SBCS_STATE_5. \heartsuit

13.6 Predicting the Behavior of the Steam-Boiler

In the previous section we have explained that failure detection was to a large extent based on a comparison between the messages received and the expected ones. For this purpose we have extended the specification SBCS_STATE by several observers, which means we have assumed that at each cycle, we record in some state variables the information needed to compute the expected messages at the next cycle. According to our explanations in Sect. 13.3, we must now specify, for each observer *obs* introduced, a corresponding *next_obs* operation. This is the aim of this section.

We have already defined the operation *next_mode* in the generic specification MODE_EVOLUTION (see Sect. 13.4) and the operation *next_status* in the generic specification STATUS_EVOLUTION (see Sect. 13.5.2). Thus what is left is the specification of the operations *next_PS_predicted*, *next_PCS_predicted*, *next_steam_predicted* and *next_level_predicted*.

As explained in Sect. 13.5, the informal requirements suggest that we should take into account some inter-dependencies when predicting values to be received at the next cycle. For instance, the water level in the steam-boiler depends on how much steam is produced, but also on how much water is poured into the steam boiler by the pumps which are open. The information provided by the water level prediction is obviously crucial to decide whether we should activate or stop some pumps. On the other hand, to predict the pump state and pump controller state messages to be received at the next cycle, we must know which pumps have been ordered to be activated or to be stopped.

From this first analysis we draw the following conclusions on how to specify the needed predictions:

- 1. In a first step we should predict the interval in which the steam output is expected to stay during the next cycle: this prediction relies only on the just received value STEAM(v) (if we trust it) or on the previously predicted values for the steam production. This is because the production of steam is expected to vary according to its maximum gradients of increase and decrease, and nothing else.
- 2. In the next step we should decide whether some pumps have to be ordered to activate or to stop. This decision, plus the knowledge about the current state of the pumps (as much as we trust it), and the predicted evolution of the steam production, should allow us to predict the evolution of the water level.
- 3. Then, on the basis of the current states of the pumps and pump controllers, together with the choice of pumps to be activated or stopped, we can predict the states of the pumps and of the pump controllers at the next cycle.

Of course all these predictions are only meaningful as long as no failure of the message transmission system has been detected (but if this is not the case the

steam-boiler control system switches to the *EmergencyStop* mode and stops, so no predictions are needed anyway). The corresponding specifications are described in the next subsections.

13.6.1 Predicting the Steam Output and the Water Level

To predict the intervals in which the steam output and the water level are expected to stay during the next cycle, we will proceed as follows (taking into account the "Additional Information" provided in [1, pp. 507–509]):

- 1. Following the analysis sketched above, when we are in the state s and have received the messages msgs, to predict the interval in which the steam output is expected to stay during the next cycle, we first should compute the *adjusted_steam* interval: this interval is either the (interval reduced to the) *received_steam* value if we can rely on it (i.e., if $PU_OK(s, msgs, SteamOutput)$ holds), or the *steam_predicted* interval (stored in the state s at the previous cycle).
- 2. Then, we use the maximum gradients of increase and decrease (i.e., U1 and U2), to predict the interval in which the steam output is expected to stay during the next cycle.
- 3. We proceed similarly for the water level: first we compute the *adjusted_level* interval, which is either the (interval reduced to the) *received_level* value if we can rely on it (i.e., if $PU_OK(s, msgs, WaterLevel)$ holds), or the *level_predicted* interval (stored in the state s at the previous cycle).
- 4. Then we should consider $broken_pumps$ (the pumps pn for which either $PU_OK(s, msgs, Pump(pn))$ does not hold or $PU_OK(s, msgs, Pump_Controller(pn))$ does not hold or both) and the $reliable_pumps$, which are not broken and are therefore known to be either *Open* or *Closed*.
- 5. At this point we must decide which pumps are ordered to activate or to stop.

However, the specific control strategy for deciding which pumps should be activated or stopped need not to be detailed in this requirements specification: this can be left to a further refinement towards an implementation of the steam-boiler control system. (Obviously the strategy should compare the adjusted_level with the recommended interval (N1, N2) and decide accordingly.)

We will therefore rely on a loosely specified *chosen_pumps* operation, for which we just impose some soundness conditions (e.g., a pump ordered to activate should be currently considered as "reliable" and *Closed*, a pump ordered to stop should be currently considered as "reliable" and *Open*).

6. Now we can compute the minimal and maximal amounts of water that will be poured into the steam-boiler during the next cycle. To compute *minimal_pumped_water*, we consider that only the pumps which are "reliable" and already *Open* will pour some water in; the *broken_pumps*, the pumps which are just ordered to activate, and the pumps which are ordered to stop are all considered not to be pouring water in. Similarly, to compute *maximal_pumped_water*, we consider that the pumps which are "reliable" and already *Open*, the pumps which are just ordered to activate, as well as all the *broken_pumps*, may pour some water in; only the "reliable" pumps just ordered to stop or already stopped are known not to be pouring any water in.

- 7. Finally, we can predict the interval in which the water level is expected to stay during the next cycle.
- 8. This prediction is the basis for deciding whether the water level risks to reach a *DangerousWaterLevel* (i.e., below *M1* or above *M2*).

Note that the intervals in which the steam output and the water level are expected to stay during the next cycle are predicted without considering the *next_status* of these devices. This is indeed necessary for the *Degraded* and *Rescue* operating modes. This leads to the following STEAM_AND_LEVEL_PREDICTION specification.

```
spec STEAM_AND_LEVEL_PREDICTION =
FAILURE_DETECTION and SET [sort PumpNumber]
then local
ops received_steam : State \times Set[R_Message] \rightarrow Value;
```

 $\begin{array}{l} adjusted_steam: State \times Set[R_Message] \rightarrow Valpair;\\ received_level: State \times Set[R_Message] \rightarrow Value;\\ adjusted_level: State \times Set[R_Message] \rightarrow Valpair;\\ broken_pumps: State \times Set[R_Message] \rightarrow Set[PumpNumber];\\ reliable_pumps: \end{array}$

 $\begin{array}{l} State \times Set[R_Message] \times PumpState \rightarrow Set[PumpNumber] \\ \forall s: State; \ msgs: Set[R_Message]; \ pn: PumpNumber; \ ps: PumpState \\ \%\% \ Axioms \ for \ STEAM: \end{array}$

- $Transmission_OK(s, msgs) \Rightarrow$ $STEAM(received_steam(s, msgs)) \in msgs$
- $adjusted_steam(s, msgs) = pair(received_steam(s, msgs), received_steam(s, msgs))$ when (Transmission_OK(s, msgs) \land PU_OK(s, msgs, SteamOutput)) else steam_predicted(s)

%% Axioms for LEVEL:

- $Transmission_OK(s, msgs) \Rightarrow$ $LEVEL(received_level(s, msgs)) \in msgs$
- adjusted_level(s, msgs) = pair(received_level(s, msgs), received_level(s, msgs)) when (Transmission_OK(s, msgs) ∧ PU_OK(s, msgs, WaterLevel)) else level_predicted(s)

%% Axioms for auxiliary pumps operations:

- $pn \ \epsilon \ broken_pumps(s, msgs) \Leftrightarrow$
 - $\neg (PU_OK(s, msgs, Pump(pn)) \land PU_OK(s, msgs, PumpController(pn)))$

• $pn \ \epsilon \ reliable_pumps(s, msgs, ps) \Leftrightarrow$ $\neg (pn \ \epsilon \ broken_pumps(s, msgs)) \land$ $PUMP_STATE(pn, ps) \ \epsilon \ msgs$

within

ops $next_steam_predicted : State \times Set[R_Message] \rightarrow Valpair;$ chosen_pumps : $State \times Set[R_Message] \times PumpState \rightarrow Set[PumpNumber];$ minimal_pumped_water, maximal_pumped_water : $State \times Set[R_Message] \rightarrow Value;$ $next_level_predicted: State \times Set[R_Message] \rightarrow Valpair$ **pred** $Dangerous WaterLevel : State \times Set[R_Message]$ %% Axioms for STEAM: $\forall s: State; msgs: Set[R_Message]; pn: PumpNumber$ • $low(next_steam_predicted(s, msqs)) =$ $max(0, low(adjusted_steam(s, msgs)) - (U2 \times dt))$ • $high(next_steam_predicted(s, msqs)) =$ $min(W, high(adjusted_steam(s, msgs)) + (U1 \times dt))$ %% Axioms for PUMPS: • $pn \ \epsilon \ chosen_pumps(s, msgs, Open) \Rightarrow$ $pn \ \epsilon \ reliable_pumps(s, msqs, Closed)$ • $pn \ \epsilon \ chosen_pumps(s, msqs, Closed) \Rightarrow$ $pn \ \epsilon \ reliable_pumps(s, msqs, Open)$ • $minimal_pumped_water(s, msqs) =$ $dt \times P \times \sharp(reliable_pumps(s, msgs, Open))$ $- chosen_pumps(s, msgs, Closed))$ • $maximal_pumped_water(s, msgs) =$ $dt \times P \times \sharp((reliable_pumps(s, msqs, Open)))$ \cup chosen_pumps(s, msgs, Open) \cup broken_pumps(s, msgs)) $- chosen_pumps(s, msgs, Closed))$ %% Axioms for LEVEL: • $low(next_level_predicted(s, msgs)) =$ $max(0, (low(adjusted_level(s, msgs)))$ $+ minimal_pumped_water(s, msqs))$ $-((dt^2 \times U1/2))$ $+ (dt \times high(adjusted_steam(s, msgs))))))$ • $high(next_level_predicted(s, msgs)) =$ $min(C, (high(adjusted_level(s, msgs)))$ $+ maximal_pumped_water(s, msgs))$ $-((dt^2 \times U2/2))$ $+ (dt \times low(adjusted_steam(s, msqs))))))$ • $DangerousWaterLevel(s, msgs) \Leftrightarrow$ $(low(next_level_predicted(s, msgs)) \leq M1) \lor$ $(M2 \leq high(next_level_predicted(s, msgs)))$
hide ops minimal_pumped_water, maximal_pumped_water
end

 \heartsuit Note the combination of *implicit hiding of auxiliary operations* declared in the *local part* and of *explicit hiding*: the operations *minimal_pumped_water* and *maximal_pumped_water* cannot be made local since their specification relies on *chosen_pumps* which must be exported. \heartsuit

13.6.2 Predicting the Pump and Pump Controller States

Specifying the predicted state of each pump at the next cycle is almost trivial. The next pump state is $Unknown_PS$ if the *next_status* of the pump is not OK, otherwise it should be Open if:

- it is *Open* now and the pump is not ordered to stop, or
- the pump is ordered to activate;

otherwise, it should be *Closed* since:

- it is *Closed* now and the pump is not ordered to activate, or
- it is ordered to stop.

This leads to the following PUMP_STATE_PREDICTION specification. This specification extends STEAM_AND_LEVEL_PREDICTION (since we rely on *chosen_pumps* for our predictions), and STATUS_EVOLUTION (which provides *next_status*) instantiated by FAILURE_DETECTION (which provides the predicate PU_OK parameter of STATUS_EVOLUTION).

end

The reasoning to predict the pump controller state is similar, but we must take into account that two cycles may be needed before a just activated pump leads to a *Flow* state (provided the pump is not stopped meanwhile). Thus, the next pump controller state is $Unknown_PCS$ if the next_status of the pump controller is not OK, or if the next_status of the corresponding pump is not OK, otherwise the predicted pump controller state value is:

- *Flow* when the pump is not ordered to stop and it is currently *Flow*, or it is currently *NoFlow* but *PCS_predicted SoonFlow*;
- NoFlow if the pump is ordered to stop, or if it is currently NoFlow and is not PCS_predicted SoonFlow and the pump is not ordered to activate;
- SoonFlow otherwise.

This leads to the following PUMP_CONTROLLER_STATE_PREDICTION specification.

```
spec PUMP_CONTROLLER_STATE_PREDICTION =
      STATUS_EVOLUTION [FAILURE_DETECTION]
      and STEAM_AND_LEVEL_PREDICTION
then op next_PCS_predicted :
            State \times Set[R\_Message] \times PumpNumber
                       \rightarrow ExtendedPumpControllerState
      \forall s: State; msgs: Set[R\_Message]; pn: PumpNumber
       • next_PCS\_predicted(s, msqs, pn) =
             Unknown_PCS when
                   \neg (next_status(s, msqs, PumpController(pn)) = OK \land
                      next\_status(s, msqs, Pump(pn)) = OK)
             else Flow when
                  (PUMP\_CONTROLLER\_STATE(pn, Flow) \in msqs \lor
                    (PUMP\_CONTROLLER\_STATE(pn, NoFlow) \in msqs \land
                      PCS\_predicted(s, pn) = SoonFlow))
                  \wedge \neg (pn \ \epsilon \ chosen\_pumps(s, msgs, Closed))
             else NoFlow when
                  (pn \ \epsilon \ chosen\_pumps(s, msgs, Closed))
                  \vee ( PUMP_CONTROLLER_STATE(pn, NoFlow) \epsilon msgs \wedge
                      \neg (PCS\_predicted(s, pn) = SoonFlow) \land
                      \neg (pn \ \epsilon \ chosen\_pumps(s, msqs, Open)))
             else SoonFlow
```

end

All our predictions are summarized in the following PU_PREDICTION specification.

```
spec PU_PREDICTION =
    PUMP_STATE_PREDICTION
    and PUMP_CONTROLLER_STATE_PREDICTION
    %{ Both specifications extend STATUS_EVOLUTION
        (instantiated by FAILURE_DETECTION)
        and STEAM_AND_LEVEL_PREDICTION }%
```

end

 \heartsuit Since the specification FAILURE_DETECTION provides the predicate PU_OK required by STATUS_EVOLUTION, we can now put pieces

together as illustrated by PU_PREDICTION. Again the 'same name, same thing' principle is essential here. \heartsuit

13.7 Specifying the Messages to Send

At this stage we are left with the specification of the messages to send at each cycle. This is easily specified, following App. C.5, and leads to the following SBCS_ANALYSIS specification.

The specification SBCS_ANALYSIS is obtained by instantiating the MODE_ EVOLUTION specification by PU_PREDICTION, and extending the result by the specification of the operation *messages_to_send*.

```
spec SBCS_ANALYSIS =
       MODE_EVOLUTION [PU_PREDICTION]
then local
       ops PumpMessages, FailureDetectionMessages :
                   State \times Set[R\_Message] \rightarrow Set[S\_Message];
            RepairedAcknowledgementMessages:
                   Set[R\_Message] \rightarrow Set[S\_Message]
       \forall s : State; msgs : Set[R\_Message]; Smsg : S\_Message
        • Smsg \in PumpMessages(s, msgs) \Leftrightarrow
             (\exists pn : PumpNumber \bullet
                    (pn \ \epsilon \ chosen\_pumps(s, msgs, Open))
                      \wedge Smsg = OPEN_PUMP(pn))
                   \lor ( pn \epsilon chosen_pumps(s, msgs, Closed)
                        \land Smsg = CLOSE\_PUMP(pn))
        • Smsg \ \epsilon \ FailureDetectionMessages(s, msgs) \Leftrightarrow
             (\exists pu : PhysicalUnit \bullet
                   Smsg = FAILURE\_DETECTION(pu) \land
                   next\_status(s, msgs, pu) = FailureWithoutAck)
        • Smsg \ \epsilon \ RepairedAcknowledgementMessages(msgs) \Leftrightarrow
             (\exists pu : PhysicalUnit \bullet
                   Smsg = REPAIRED\_ACKNOWLEDGEMENT(pu) \land
                   next\_status(s, msgs, pu) = FailureWithAck)
       within
       op messages_to_send : State \times Set[R_Message] \rightarrow Set[S_Message]
       \forall s: State; msqs: Set[R\_Message]
        • messages\_to\_send(s, msgs) =
             (PumpMessages(s, msgs) \cup
              FailureDetectionMessages(s, msgs) \cup
              RepairedAcknowledgementMessages(msgs))
              + MODE(next\_mode(s, msgs))
```

end

 \heartsuit We rely again here on *auxiliary operations* declared in the *local part*, and their axiomatization is fairly easy using *existential quantifiers*. \heartsuit

13.8 The Steam-Boiler Control System Specification

According to our work plan detailed in Sect. 13.3, we have already specified the main parts of our case study. First, let us display a basic (flat) specification equivalent to SBCS_STATE_5 and where all the state observers are listed together.

\mathbf{spec}	$SBCS_STATE =$
	Preliminary
then	sort State
	free type $Status ::= OK \mid Failure WithoutAck \mid Failure WithAck$
	free type <i>ExtendedPumpState</i> ::= <i>sort PumpState</i> <i>Unknown_PS</i>
	free type
	<i>ExtendedPumpControllerState</i> ::= sort PumpControllerState
	SoonFlow Unknown_PCS
	free type $Valpair ::= pair(low : Value; high : Value)$
	ops $mode: State \rightarrow Mode;$
	$numSTOP: State \rightarrow Nat;$
	$status: State \times PhysicalUnit \rightarrow Status;$
	$PS_predicted : State \times PumpNumber$
	\rightarrow ExtendedPumpState;
	$PCS_predicted : State \times PumpNumber$
	\rightarrow ExtendedPumpControllerState;
	$steam_predicted$, $level_predicted : State \rightarrow Valpair$

end

We are now ready to provide the specification of the steam-boiler control system, considered as a labeled transition system. We leave partly unspecified the initial state *init*, since in our specification this state represents the state immediately following the receipt of the *PHYSICAL_UNITS_READY* message. Hence intuitively the omitted axioms should take into account the messages sent and received during the initialization phase (at least at the end of it). It is therefore better to leave open for now the value of most observers on *init*, and to note that this would have to be taken care of when specifying the initialization phase. The value of *mode(init)* is specified according to the end of App. C.4.1.

spec STEAM_BOILER_CONTROL_SYSTEM =
 SBCS_ANALYSIS
then op init : State
 pred is_step : State × Set[R_Message] × Set[S_Message] × State

%% Specification of the initial state *init*: • $mode(init) = Normal \lor mode(init) = Degraded$ %% Specification of *is_step*: $\forall s, s' : State; msgs : Set[R_Message]; Smsq : Set[S_Message]$ • $is_step(s, msqs, Smsq, s') \Leftrightarrow$ $mode(s') = next_mode(s, msqs) \land$ $numSTOP(s') = next_numSTOP(s, msqs) \land$ $(\forall pu : PhysicalUnit \bullet$ $status(s', pu) = next_status(s, msqs, pu)) \land$ $(\forall pn : PumpNumber \bullet$ $PS_predicted(s', pn) = next_PS_predicted(s, msqs, pn) \land$ $PCS_predicted(s', pn) = next_PCS_predicted(s, msgs, pn)) \land$ $steam_predicted(s') = next_steam_predicted(s, msgs) \land$ $level_predicted(s') = next_level_predicted(s, msgs) \land$ $Smsg = messages_to_send(s, msgs)$ then %% Specification of the reachable states: free { pred reach : State $\forall s, s' : State; msgs : Set[R_Message]; Smsg : Set[S_Message]$ • reach(init)• $reach(s) \land is_step(s, msgs, Smsg, s') \Rightarrow reach(s')$ }

end

13.9 Validation of the CASL Requirements Specification

Once the formalization of the informal requirements is completed, we must now face the following question: is our formal specification adequate? This is a difficult question to answer since there is no formal way to establish the adequacy of a formal specification w.r.t. informal requirements, i.e., we cannot *prove* this adequacy. However, we can try to *test* it, by performing various 'experiments'. When these experiments are successful, our confidence in the formal specification is increased. If some experiment fails, then we can inspect the specification and try to understand the causes of the failure, possibly detecting some flaw in the specification.

We will base our validation process on theorem proving, i.e., we will check that some formulas are logical consequences of our requirements specification STEAM_BOILER_CONTROL_SYSTEM. For this purpose we use the tools described in Chap. 11. During this validation process we can consider two kinds of proof obligations:

1. We can inspect the text of the specification and derive from this inspection some formulas that are expected to be logical consequences of our specification. This can be considered as a kind of internal validation of the formal specification. 2. We can check that some expected properties inferred from the informal requirements are logical consequences of our specification (external validation). To do this, we must first reanalyze the informal specification, state some expected properties, translate them into formulas, and then attempt to prove that these formulas are logical consequences of our specification. This task is not easy, since in general one has the feeling that all expected properties were already detected and included in the axioms during the formalization process.

The application of these principles to the requirements specification of the steam-boiler control system leads to various proofs. Below we give only a few illustrative examples.

For instance, let us consider the specification of *next_mode* in MODE_EVOLUTION: it is advisable to prove that all the cases considered in the specification of *next_mode* are mutually exclusive, and that their disjunction is equivalent to true. This is a typical example of internal validation of the specification, since we just consider the text of the specification to decide which proof attempt will be performed, without considering the informal requirements again. We do not spell out the corresponding proofs here, but the reader can easily check that indeed the operation *next_mode* is well-defined (i.e., all cases are mutually exclusive and their disjunction is equivalent to true). In the same spirit we can prove that the same pump cannot simultaneously be ordered to activate and to stop, that we never resignal a failure which has already been signaled, that as long as the operating mode is not set to *EmergencyStop* the water level is safe, etc.

Let us now consider an example of external validation. According to our understanding of failure detection (see Sect. 13.5), if we have detected a failure of some physical unit pu (so PU_OK does not hold for pu), then the status of this physical unit should not be set to OK. The corresponding proof obligation reads as follows:

STEAM_BOILER_CONTROL_SYSTEM \models $\forall s : State; msgs : Set[R_Message]; pu : PhysicalUnit$ • Transmission_OK(s, msgs) $\land \neg PU_OK(s, msgs, pu)$ $\land reach(s) \Rightarrow \neg (next_status(s, msgs, pu) = OK)$

However here we are unable to discharge this proof obligation. A careful analysis of the proof attempt shows that the proof fails since it could be the case that, simultaneously with the receipt of a repaired message for the physical unit pu, we nevertheless detect again a failure of the same unit. From this analysis we conclude that the following axiom in STATUS_EVOLUTION is not adequate:

status(s, pu) = FailureWithAck ∧ REPAIRED(pu) is_in msgs
 ⇒ next_status(s, msgs, pu) = OK

This means we must fix the STATUS_EVOLUTION specification and replace the above axiom by:

status(s, pu) = FailureWithAck ∧ REPAIRED(pu) is_in msgs
 ⇒ next_status(s, msgs, pu) = OK when PU_OK(s, msgs, pu) else FailureWithoutAck

Once the specification STATUS_EVOLUTION is modified as explained above, we can prove that the expected property holds.

To conclude, the reader should keep in mind that the validation of the specification is a very important task that deserves some serious attention. In this section we have only briefly illustrated some typical proof attempts that would naturally arise when validating the STEAM_BOILER_CONTROL_SYSTEM specification, and obviously many other proof attempts are required to reach a stage where we can trust our requirements specification of the steam-boiler control system.

13.10 Designing the Architecture

We now have a validated requirements specification STEAM_BOILER_CONTROL_ SYSTEM of the steam-boiler control system. The next step is to refine it into an architectural specification, thereby prescribing the intended architecture of the implementation of the steam-boiler control system. Indeed, the explanations given in Sect. 13.3 suggest the following rather obvious architecture for the steam-boiler control system:

```
arch spec Arch_SBcs =

units P: Value \rightarrow Preliminary;

S: Preliminary \rightarrow SBcs_State;

A: SBcs_State \rightarrow SBcs_Analysis;

C: SBcs_Analysis \rightarrow Steam_Boiler_Control_System

result \lambda V: Value • C[A[S[P[V]]]]

end
```

Note that we decide to describe the implementation of the steam-boiler control system as an *open system*, relying on an external component V implementing VALUE. This is consistent with our explanations in Sect. 13.2: choosing a specific implementation of VALUE is obviously orthogonal to designing the implementation of the steam-boiler control system. This means in particular that the component V implementing VALUE will encapsulate the chosen representation of natural numbers and values, together with operations and predicates operating on them.

 \heartsuit As illustrated by ARCH_SBCS, the intended architecture of the steam-boiler control system is easily described by an *architectural specification*. Then we can proceed with four separate implementation tasks, which are independent of each other. \heartsuit

In a next step, we can refine the specification VALUE \rightarrow PRELIMINARY of the component P into the following architectural specification.

```
arch spec Arch_Preliminary =

units SET : { sort Elem} × NAT \rightarrow SET [sort Elem];

B : BASICS;

MS : MESSAGES_SENT given B;

MR : VALUE \rightarrow MESSAGES_RECEIVED given B;

CST : VALUE \rightarrow SBCS_CONSTANTS

result \lambda V : VALUE • SET [MS fit Elem \mapsto S_Message] [V]

and SET [MR [V] fit Elem \mapsto R_Message] [V]

and CST [V]

end
```

end

Here we decide to implement (generic) sets in a component SET, reused both for sets of messages received and sets of messages sent. Since the implementation of natural numbers is provided by the (external) component V, we use V for the second argument of the generic component SET in the result unit term.

 \heartsuit Note how the generic specification with imports SET is transposed into a specification of a generic component SET. Note also, for the component MR, the use of a specification of a generic component extending a given unit. \heartsuit

The specification of the components C and S of ARCH_SBCS are simple enough that they do not need to be further architecturally refined. The specification of the component S (which implements the states of the steam-boiler control system) can be refined into the following specification UNIT_SBCS_ STATE, which provides a concrete implementation of states as a record of all the observable values.

from BASIC/STRUCTUREDDATATYPES get TOTALMAP

spec	$SBCS_STATE_IMPL =$
	Preliminary
\mathbf{then}	free type $Status ::= OK \mid Failure WithoutAck \mid Failure WithAck$
	free type ExtendedPumpState ::= sort PumpState Unknown_PS
	free type ExtendedPumpControllerState ::=
	$sort PumpControllerState \mid SoonFlow \mid Unknown_PCS$
	free type $Valpair ::= pair(low : Value; high : Value)$
then	TOTALMAP [BASICS fit $S \mapsto PhysicalUnit$] [sort Status]
and	TOTALMAP [BASICS fit $S \mapsto PumpNumber$] [sort ExtendedPumpState]
and	TOTALMAP [BASICS fit $S \mapsto PumpNumber$]
	[sort $ExtendedPumpControllerState]$

```
then free type State ::= mk\_state(
                 mode : Mode:
                 numSTOP : Nat;
                 status : TotalMap[PhysicalUnit, Status];
                 PS\_predicted :
                   TotalMap[PumpNumber, ExtendedPumpState];
                 PCS\_predicted:
                   TotalMap[PumpNumber, ExtendedPumpControllerState];
                 steam_predicted, level_predicted : Valpair )
      ops status(s : State; pu : PhysicalUnit) : Status
                 = lookup(pu, status(s));
           PS\_predicted(s:State; pn:PumpNumber):ExtendedPumpState
                 = lookup(pn, PS\_predicted(s));
           PCS\_predicted(s:State; pn:PumpNumber)
                 : ExtendedPumpControllerState
                 = lookup(pn, PCS\_predicted(s))
```

end

```
unit spec UNIT_SBCS_STATE = 
 PRELIMINARY \rightarrow SBCS_STATE_IMPL
```

 \heartsuit During the formalization process it was convenient to rely on a *loose* specification of states. At the design stage, this loose specification is refined into a specification where state variables are now explicit. \heartsuit

The specification SBCS_STATE \rightarrow SBCS_ANALYSIS of the component A of ARCH_SBCS can be refined into the following architectural specification:

```
arch spec Arch_Analysis =

units FD : SBCS_STATE \rightarrow FAILURE_DETECTION;

PR : FAILURE_DETECTION \rightarrow PU_PREDICTION;

ME : PU_PREDICTION \rightarrow MODE_EVOLUTION [PU_PREDICTION];

MTS : MODE_EVOLUTION [PU_PREDICTION] \rightarrow SBCS_ANALYSIS

result \lambda S : SBCS_STATE • MTS [ME [PR [FD [S]]]]

end
```

In the above architectural specification ARCH_ANALYSIS, the component FD provides an implementation of failure detection, the component PR an implementation of the predicted state variables for the next cycle, the component ME provides an implementation of next_mode (and of next_numSTOP), and the component MTS provides an implementation of messages_to_send.

The specifications of the components ME and MTS are simple enough to be directly implemented. The specifications of the components FD and PRcan be refined as follows.

```
arch spec Arch_FAILURE_DETECTION =
units MTSF : SBCS_STATE
                     \rightarrow Message_Transmission_System_Failure;
      PF
             : SBCS_STATE \rightarrow PUMP_FAILURE:
             : SBCS_STATE \rightarrow PUMP_CONTROLLER_FAILURE;
      PCF
      SF
             : SBCS_STATE \rightarrow STEAM_FAILURE;
             : SBCS_STATE \rightarrow Level_Failure:
      LF
      PU
             : Message_Transmission_System_Failure
               × Pump_Failure × Pump_Controller_Failure
               × Steam Failure × Level Failure
                     \rightarrow Failure_Detection
result \lambda S : SBCS_STATE •
            PU [MTSF[S]] [PF[S]] [PCF[S]] [SF[S]] [LF[S]]
            hide Pump_OK, Pump_Controller_OK, Steam_OK, Level_OK
```

end

The above architectural specification ARCH_FAILURE_DETECTION refines the specification SBCS_STATE \rightarrow FAILURE_DETECTION of the component FDin ARCH_ANALYSIS and introduces a component for each kind of failure detection. Then the component PU implements PU_OK , and in the result unit expression we hide the auxiliary predicates provided by the components PF, PCF, SF, and LF.¹⁰

We refine the specification FAILURE_DETECTION \rightarrow PU_PREDICTION of the component PR of the architectural specification ARCH_ANALYSIS as follows:

```
arch spec Arch_Prediction =

units SE : FAILURE_DETECTION \rightarrow

STATUS_EVOLUTION [FAILURE_DETECTION];

SLP : FAILURE_DETECTION \rightarrow STEAM_AND_LEVEL_PREDICTION;

PP : STATUS_EVOLUTION [FAILURE_DETECTION]

\times STEAM_AND_LEVEL_PREDICTION

\rightarrow PUMP_STATE_PREDICTION;

PCP : STATUS_EVOLUTION [FAILURE_DETECTION]

\times STEAM_AND_LEVEL_PREDICTION

\rightarrow PUMP_CONTROLLER_STATE_PREDICTION

result \lambda FD : FAILURE_DETECTION •

local SEFD = SE [FD]; SLPFD = SLP [FD] within

PP [SEFD] [SLPFD] and PCP [SEFD] [SLPFD]
```

end

¹⁰ These auxiliary predicates are already hidden in the specification FAILURE_ DETECTION. However, remember that in the specification of a generic component, the target specification is always an implicit extension of the argument specifications. This is why it is necessary to hide the auxiliary predicates at the level of the result unit expression.

190 13 Case Study: The Steam-Boiler Control System

In the above architectural specification, the component SE provides an implementation of *next_status*. The component SLP provides an implementation of *next_steam_predicted*, *next_level_predicted*, *chosen_pumps*, and *Dangerous-WaterLevel*. The component PP provides an implementation of *next_PS_predicted*, and the component PCP provides an implementation of *next_PCS_predicted*.

We are now left with specifications of components that are simple enough to be directly implemented, and this concludes our case study.

CASL Quick Reference

This appendix provides an overview of the (concrete) syntax of each part of CASL.

Basic specifications

- declarations, definitions:
 - sorts, subsorts
 - functions: total, partial
 - constants: total, partial
 - predicates
 - datatypes
 - sort generation constraints
- variables, axioms
 - formulas
 - terms
- symbols
- comments
- annotations

$Architectural\ specifications$

- named architectures, units
- architectural specifications
- unit specifications
- unit declarations, definitions
- unit expressions, terms

Structured specifications

- specification structure
 - translation
 - hiding, revealing
 - union, extension
 - free extension, initiality
 - hiding local symbols
 - reference
 - instantiation
- named, generic specifications
 - fitting arguments
- named, generic views
 - fitting views
- symbol lists, maps

Libraries

- named libraries
- downloadings
- library names, versions

A.1 Basic Specifications

;	list of items (';' optional)
sorts	sort declarations and definitions
$ops \dots$	operation declarations and definitions
preds	predicate declarations and definitions
types	datatype declarations and definitions
generated $\{\ldots\}$	sort generation constraint
vars	global variable declarations
$\forall \dots \bullet F_1 \dots \bullet F_n$	universally-quantified list of axioms
• $F_1 \ldots \bullet F_n$	unquantified list of axioms

A.1.1 Declarations and Definitions

Sort Declarations and Definitions

\mathbf{sort}	S	sort declaration
sorts	s_1,\ldots,s_n	sorts declaration
\mathbf{sorts}	s < s'	subsort declaration
sorts	$s_1, \ldots, s_n < s'$	subsorts and supersort declaration
sorts	$s < s_1; \ldots; s < s_n$	subsort and supersorts declaration
\mathbf{sorts}	$s_1 = \ldots = s_n$	isomorphic sorts declaration
\mathbf{sort}	$s = \{v : s' \bullet F\}$	subsort definition

Function Declarations and Definitions

op $f: s_1 \times \cdots \times s_n \to s$
op $f: s_1 \times \cdots \times s_n \to ? s$
op $f: s \times s \to s, assoc$
op $f: s \times s \to s', comm$
op $f: s \times s \to s, idem$
op $f: s \times s \to s, unit T$
op $f: s \times s \to s, \dots, \dots$
$\mathbf{ops} f_1, \ldots, f_n : \ldots$
op $f(v_1:s_1;\ldots;v_n:s_n):s=T$
op $f(v_1 : s_1; \ldots; v_n : s_n) :?s = T$
op $f(\ldots v_{i_1}, \ldots, v_{i_m} : s_i \ldots) \ldots$
ops;

total function declaration partial function declaration associative binary function commutative binary function idempotent binary function unit term for binary function multiple function attributes functions declaration total function definition partial function definition abbreviated arguments multiple declarations/definitions

Constant Declarations and Definitions

c:s	constant declaration
c:?s	partial constant declaration
$c_1,\ldots,c_n:s$	constants declaration
c: s = T	constant definition
c:?s=T	partial constant definition
;	multiple declarations/definitions
	$c: s$ $c:?s$ $c_1, \dots, c_n: s$ $c: s = T$ $c:?s = T$ $\dots; \dots$

Predicate Declarations and Definitions

 $\begin{array}{lll} \mathbf{pred} & p:s_1 \times \cdots \times s_n \\ \mathbf{pred} & p:() \\ \mathbf{preds} & p_1, \ldots, p_n: \ldots \\ \mathbf{pred} & p(v_1:s_1; \ \ldots; \ v_n:s_n) \Leftrightarrow F \\ \mathbf{pred} & p \Leftrightarrow F \\ \mathbf{pred} & p(\ldots v_{i_1}, \ldots, v_{i_m}:s_i \ldots) \ldots \\ \mathbf{preds} & \ldots; & \ldots \end{array}$

predicate declaration constant predicate declaration predicates declaration predicate definition constant predicate definition abbreviated arguments multiple declarations/definitions

Datatype Declarations

$\mathbf{type} \ s ::= A$	datatype declaration with alternatives
$\mathbf{types} \ s_1 ::= A_1;$	multi-sorted datatype declaration
;	
$s_n ::= A_n$	
generated types	generated datatype declaration
free types	free datatype declaration

Alternatives (A)

total constructor function partial constructor function total constructor and selector functions total constructor, partial selector functions abbreviated selectors constant constructor value subsort subsort multiple alternatives

Sort Generation Constraints

generated $\{ \text{ sorts } \dots \}$	generated sorts
$\operatorname{ops}\ldots$	generating operations
$preds \dots$	
$\mathbf{types} \dots$	generated sorts
}	and generating constructors

A.1.2 Variables and Axioms

```
\begin{array}{l} \mathbf{var} \ v: s \\ \mathbf{vars} \ v_1: s_1; \ \dots; \ v_n: s_n \\ \mathbf{vars} \ \dots; \ v_1, \dots, v_n: s_n \dots \\ \mathbf{vars} \ \dots; \ \dots \\ \forall v: s \ \bullet \ F_1 \dots \ \bullet \ F_n \\ \forall v_1, \dots, v_n: s \ \bullet \ \dots \\ \bullet \ F_1 \dots \ \bullet \ F_n \end{array}
```

Formulas (F)

$$\forall \dots \bullet F \\ \exists \dots \bullet F \\ \exists ! \dots \bullet F \\ F_1 \land \dots \land F_n \\ F_1 \lor \dots \lor F_n \\ F \Rightarrow F' \\ F' if F \\ F \Leftrightarrow F' \\ \neg F \\ true \\ false \\ p(T_1, \dots, T_n) \\ t_0 \ T_1 \ t_1 \dots T_n \ t_n \\ q \\ T = T' \\ T \stackrel{e}{=} T' \\ def \ T \\ T \in s$$

Terms (T)

$$f(T_1, \dots, T_n)$$

$$t_0 T_1 t_1 \dots T_n t_n$$

$$t_0 T_1, \dots, T_n t_1$$

$$c$$

$$v$$

$$T: s$$

$$T as s$$

$$T when F else T'$$

global variable declaration global variables declaration abbreviated variables declaration multiple global variable declarations universally-quantified list of axioms abbreviated quantifications multiple quantifications unquantified list of axioms

universal quantification on formula existential quantification unique-existential quantification conjunction disjunction implication reverse implication equivalence negation truth falsity predicate application mixfix predicate application constant predicate ordinary (strong) equality existential equality definedness subsort membership

application mixfix application literal syntax constant variable sorted term projection to subsort conditional choice

A.1.3 Symbols

Character set: ASCII (with optional use of ISO Latin-1).

Key Words and Signs

Reserved key words (always lowercase):

and arch as axiom axioms closed def else end exists false fit forall free from generated get given hide if in lambda library local not op ops pred preds result reveal sort sorts spec then to true type types unit units var vars version view when with within

Reserved key signs:

: :? ::= = => <=> \neg . · | |-> \/ /\

Unreserved key signs:

< * \times -> ? ! [] { }

Key words and signs representing mathematical symbols:

forall	exists	exists!	not	in	lambda	=e=	->	=>	<=>		•	->	\land	$\setminus/$
\forall	Ξ	∃!	_	\in	λ	$\stackrel{e}{=}$	\rightarrow	\Rightarrow	\Leftrightarrow	•	•	\mapsto	\wedge	\vee

Identifiers

Identifiers for sorts and variables are simple words (other than reserved words) possibly containing digits, primes, and *single* underscores:

Elem Y_1 Z2' A_Rather_Long_Identifier

Sort identifiers can also be compound:

List[Int] Map[Index, Elem]

Identifier for operations and predicates can moreover be sequences of (unreserved) signs, with any brackets [] { } balanced:

 $+ - * / \setminus \& = <>[] \{ \} ! ? : . \$ @ \# ^ ~ [: × ÷ \pounds © ± \P \S ^ 1 ? ^ 3 \cdot \not e^{\circ} \neg \mu |$

or single decimal digits 1 2 3 4 5 6 7 8 9 0, or single quoted characters 'c'.

The signs (); ,' "% are *not* allowed in identifiers, nor are the ISO Latin-1 signs for general currency, yen, broken vertical bar, registered trade mark, masculine and feminine ordinals, left and right angle quotes, fractions, soft hyphen, acute accent, cedilla, macron, and umlaut.

Operation and predicate identifiers can also be compound:

 $order[_] < _]$

Function and predicate identifiers can also be infixes, prefixes, postfixes, and general mixfixes, formed from words and/or sequences of signs separated by *double* underscores (indicating the positions of the arguments), with any brackets [] { } balanced:

__++__ ||__|| {[_]} push_onto__ _.select1

Invisible mixfix identifiers (such as ____) with two or more arguments are allowed. (Subsort embeddings give the effect of invisible unary functions.)

An operation, or predicate identifier can be compound, with a list of identifiers appended to its final token.

Literal Strings and Numbers

"this is a string" 42 3.14159 1E-9 27.3e6

Library Identifiers

Names of libraries are either paths, e.g.:

BASIC/NUMBERS BASIC/ALGEBRA_II

or URLs formed from A...ZA...Z0...9-@.&+!*"'(),~ and hexadecimal codes %xx, and prefixed by HTTP://, FTP://, or FILE:///.

Version numbers of libraries are hierarchical: 0, 0.999, 1, 1.0, 1.0.2.

A.1.4 Comments

%% This is a comment at the end of a line. . .

- $\dots \%$ This is an in-line comment $\}\% \dots$
- ...% { This a comment that might take several lines }%

%[This is for commenting-out text %{ including other kinds of comment }%]%

A.1.5 Annotations

A label is of the form %(text)%.

An end-of-line annotation is of the general form %word ... with a space following the word.

A possibly multi-line annotation is of the general form %word(...)% with no space preceding the '('.

A.2 Structured Specifications

A.2.1 Specifications (SP)

symbol translation
hiding listed symbols
revealing/translating listed symbols
union
extension
free or initial
hiding of local symbols
self-contained
reference to named specification
instantiation of generic specification

A.2.2 Named and Generic Specifications

spec $SN = SP$ end	named specification (end optional)
spec $SN[SP_1][SP_n] = SP$ end	generic specification (end optional)
spec $SN[SP_1][SP_n]$	generic specification
given $SP''_1, \ldots, SP''_m = SP$ end	with imports (end optional)

Fitting Arguments (FA)

SP fit SM	fitting by symbol map
SP	implicit fitting
FV	fitting view

A.2.3 Named and Generic Views

view	VN: SP to $SP' = SM$ end	named view (end optional)
view	$VN[SP_1]\dots[SP_n]$	generic view
	: SP to $SP' = SM$ end	(end optional)
view	$VN[SP_1]\dots[SP_n]$	generic view
	given SP''_1, \ldots, SP''_m	with imports
	: SP to $SP' = SM$ end	(end optional)

Fitting Views (FV)

view	VN	reference to named view
view	$VN[FA_1]\dots[FA_n]$	instantiation of generic view

A.2.4 Symbol Lists (SL) and Maps (SM)

SY_1, \ldots, SY_n	lists (maybe with sorts , ops , preds)
$SY_1 \mapsto SY'_1, \dots, SY_n \mapsto SY'_n$	maps (maybe with $sorts$, ops , $preds$)
\ldots, SY_i, \ldots	in a map, abbreviates $\ldots, SY_i \mapsto SY_i, \ldots$

A.3 Architectural Specifications

A.3.1 Named Architectures and Units

arch spec $ASN = ASP$ end	named arch. spec. (end optional)
unit spec $SN = USP$ end	named unit spec. (end optional)

A.3.2 Architectural Specifications (ASP)

ASN				${\rm arch.}$	spec.	name
\mathbf{units}	$UD_1; \ldots;$	UD_n result	UE	basic	arch.	spec.

A.3.3 Unit Specifications (USP)

SP	unit specification
$SP_1 \times \ldots \times SP_n \to SP$	generic-unit specification
closed USP	self-contained
arch spec ASP	models of arch. spec.

A.3.4 Unit Declarations and Definitions (UD)

UN: USP	unit declaration
$UN: USP$ given UT_1, \ldots, UT_n	importing units
UN = UE	unit definition

A.3.5 Unit Expressions (UE)

UT				unit term
$\lambda UN_1 : SP_1; \ldots;$	$UN_n:SP_n$	•	UT	unit composition

A.3.6 Unit Terms (UT)

UT with SM	symbol translation
UT hide SL	hiding listed symbols
UT reveal SM	revealing/translating listed symbols
UT_1 and and UT_n	amalgamation
local $UD_1; \ldots; UD_n$ within UT	local units
UN	unit name
$UN[UT_1]\dots[UT_n]$	generic-unit application
$UN[UT_1 \text{ fit } SM_1][UT_n \text{ fit } SM_n]$	with fitting by symbol maps

A.4 Libraries

library LN	named library of downloadings,
	specifications, views

A.4.1 Downloadings

from LN **get** IN_1, \ldots, IN_n **end** downloads listed items **from** LN **get** $\ldots IN \mapsto IN' \ldots$ **end** renames downloaded items

A.4.2 Library Names (LN)

Basic/Numbers	greatest version registered
BASIC/ALGEBRA_II version 0.999	specified version registered
НТТР://	greatest version unregistered
HTTP:// version $1.0.2$	specified version unregistered

Points to Bear in Mind

B.1 Introduction

•	CoFI aims at establishing a wide consensus.	4
•	The focus of CoFI is on <i>algebraic</i> techniques	5
•	CoFI has already achieved its main aims	5
•	CoFI is an open, voluntary initiative	6
•	CoFI has received funding as an ESPRIT Working Group, and is	
	sponsored by IFIP WG 1.3.	6
•	New participants are welcome!	7
•	CASL has been designed as a general-purpose algebraic specification	
	language, subsuming many existing languages.	7
•	CASL is at the center of a <i>family</i> of languages	8
•	CASL itself has several major <i>parts</i>	9

B.2 Underlying Concepts

•	CASL is based on standard concepts of algebraic specification	11
•	A basic specification declares symbols, and gives axioms and	
	constraints	11
•	The semantics of a basic specification is a signature and a class of	
	models	12
•	CASL specifications may declare sorts, subsorts, operations, and	
	predicates.	12
•	Sorts are interpreted as carrier sets	12
•	Subsorts declarations are interpreted as embeddings	13
•	Operations may be declared as total or partial	13
•	Predicates are different from boolean-valued operations	13
•	Operation symbols and predicate symbols may be overloaded	14
•	Axioms are formulas of first-order logic	14

•	Sort generation constraints eliminate 'junk' from specific carrier sets.	15
•	The semantics of a structured specification is simply a signature	
	and a class of models.	16
•	A translation merely renames symbols	17
•	Hiding symbols removes parts of models	17
•	Union of specifications identifies common symbols	17
•	Extension of specifications identifies common symbols too	18
•	Free specifications restrict models to being free, with initiality as a	
	special case	18
•	Generic specifications have parameters, and have to be instantiated	
	when referenced	18
•	The semantics of an architectural specification reflects its modular	
	structure	19
•	Architectural specifications involve the notions of persistent	
	function and conservative extension	19
•	The semantics of a library of specifications is a mapping from the	
	names of the specifications to their semantics.	20

B.3 Getting Started

•	Simple specifications may be written in CASL essentially as in many	
	other algebraic specification languages.	23
•	CASL provides also useful abbreviations	23
•	CASL allows loose, generated and free specifications.	24
•	CASL syntax for declarations and axioms involves familiar notation,	
	and is mostly self-explanatory.	24
•	Specifications can easily be extended by new declarations and axioms.	25
•	In simple cases, an operation (or a predicate) symbol may be	
	declared and its intended interpretation defined at the same time	26
•	Symbols may be conveniently displayed as usual mathematical	
	symbols by means of %display annotations	27
•	The %implies annotation is used to indicate that some axioms are	
	supposedly redundant, being consequences of others	28
•	Attributes may be used to abbreviate axioms for associativity,	
	commutativity, idempotence, and unit properties	29
•	Genericity of specifications can be made explicit using parameters	29
•	References to generic specifications always instantiate the parameters.	30
•	Datatype declarations may be used to abbreviate declarations of	
	sorts and constructors	32
•	Loose datatype declarations are appropriate when further	
	constructors may be added in extensions	32
•	Sorts may be specified as generated by their constructors	33
•	Generated specifications are in general loose	34
•	Generated specifications need not be loose	35

•	Generated types may need to be declared together	36
•	Free specifications provide initial semantics and avoid the need for	
	explicit negation.	36
•	Free datatype declarations are particularly convenient for defining	
	enumerated datatypes	37
•	Free specifications can also be used when the constructors are	
	related by some axioms.	37
•	Predicates hold minimally in models of free specifications	38
•	Operations and predicates may be safely defined by induction on	
	the constructors of a free datatype declaration.	38
•	More care may be needed when defining operations or predicates	
	on free datatypes when there are axioms relating the constructors	39
•	Generic specifications often involve free extensions of (loose)	
	parameters.	40
•	Loose extensions of free specifications can avoid overspecification	41
•	Datatypes with observer operations or predicates can be specified	
	as generated instead of free.	42
•	The %def annotation is useful to indicate that some operations or	
	predicates are uniquely defined.	43
•	Operations can be defined by axioms involving observer operations,	
	instead of inductively on constructors.	44
•	Sorts declared in free specifications are not necessarily generated	
	by their constructors.	44

B.4 Partial Functions

•	Partial functions arise naturally.	47
•	Partial functions are declared differently from total functions	47
•	Terms containing partial functions may be undefined, i.e., they may	
	fail to denote any value.	48
•	Functions, even total ones, propagate undefinedness	48
•	Predicates do not hold on undefined arguments	48
•	Equations hold when both terms are undefined	48
•	Special care is needed in specifications involving partial functions	49
•	The definedness of a term can be checked or asserted	50
•	The domains of definition of partial functions can be specified exactly.	50
•	Loosely specified domains of definition may be useful	51
•	Domains of definition can be specified more or less explicitly	51
•	Partial functions are minimally defined by default in free	
	specifications	53
•	Selectors can be specified concisely in datatype declarations, and	
	are usually partial.	54
•	Selectors are usually total when there is only one constructor	54
•	Constructors may be partial.	54

•	Existential equality requires the definedness of both terms as well	
	as their equality.	55

B.5 Subsorting

•	Subsorts and supersorts are often useful in CASL specifications Subsort declarations directly express relationships between carrier	57
	sets	57
•	Operations declared on a sort are automatically inherited by its	
	subsorts	58
•	Inheritance applies also for subsorts that are declared afterwards	59
•	Subsort membership can be checked or asserted	59
•	Datatype declarations can involve subsort declarations	59
•	Subsorts may also arise as classifications of previously specified	
	values, and their values can be explicitly defined	60
•	It may be useful to redeclare previously defined operations, using	
	the new subsorts introduced	61
•	A subsort may correspond to the definition domain of a partial	
	function	62
•	Using subsorts may avoid the need for partial functions	62
•	Casting a term from a supersort to a subsort is explicit and the	
	value of the cast may be undefined.	63
•	Supersorts may be useful when generalizing previously specified sorts. Supersorts may also be used for extending the intended values by	64
	new values representing errors or exceptions	65

B.6 Structuring Specifications

•	Large and complex specifications are easily built out of simpler ones	
	by means of (a small number of) specification-building operations	67
•	Union and extension can be used to structure specifications	67
•	Specifications may combine parts with loose, generated, and free	
	interpretations	68
•	Renaming may be used to avoid unintended name clashes, or to	
	adjust names of sorts and change notations for operations and	
	predicates	69
•	When combining specifications, origins of symbols can be indicated	71
•	Auxiliary symbols used in structured specifications can be hidden	71
•	Auxiliary symbols can be made local when they do not need to be	
	exported	73
•	Care is needed with local sort declarations	74
•	Naming a specification allows its reuse	75

B.7 Generic Specifications

•	Making a specification generic (when appropriate) improves its	77
•	Parameters are arbitrary specifications.	78
•	The argument specification of an instantiation must provide	
•	symbols corresponding to those required by the parameter The argument specification of an instantiation must ensure that the	78
•	properties required by the parameter hold	79
•	specification and the body of the instantiated generic specification In instantiations, the fitting of parameter symbols to identical	80
	argument symbols can be left implicit	80
•	left implicit.	80
•	implicit too, and can imply fitting of sorts.	81
•	I he intended fitting of the parameter symbols to the argument	01
	A generic specification may have more than one parameter	82
•	Instantiation of generic specifications with several parameters is	02
	similar to the case of just one parameter.	82
•	Composition of generic specifications is expressed using instantiation. Compound sorts introduced by a generic specification get	84
	automatically renamed on instantiation, which avoids name clashes	85
•	Compound symbols can also be used for operations and predicates Parameters should be distinguished from references to fixed	87
•	specifications that are not intended to be instantiated Argument specifications are always implicitly regarded as extension	88
	of the imports.	89
•	Imports are also useful to prevent ill-formed instantiations	89
•	In generic specifications, auxiliary required specifications should be	00
•	Views are named fitting maps, and can be defined along with	90
•	specifications	90 91
	0	

B.8 Specifying the Architecture of Implementations

•	Architectural specifications impose structure on implementations,	
	whereas specification-building operations only structure the text of	
	specifications.	93

•	An architectural specification consists of a list of unit declarations,
	how they are to be combined 96
•	There can be several distinct architectural choices for the same
-	requirements specification
•	Each unit declaration listed in an architectural specification
	corresponds to a separate implementation task
•	A unit can be implemented only if its specification is a conservative
	extension of the specifications of its given units
•	Genericity of components can be made explicit in architectural
	specifications
•	A generic component may be applied to an argument richer than
	required by its specification
•	Specifications of components can be named for further reuse 102
•	Both named and unnamed specifications can be used to specify
	components
•	specifications of generic components should not be comused with
	A generic component may be applied more than once in the same
•	architectural specification 103
•	Several applications of the same generic component is different from
	applications of several generic components with similar specifications. 104
•	Generic components may have more than one argument
•	Open systems can be described by architectural specifications using
	generic unit expressions in the result part
•	When components are to be combined, it is best to check that any
	shared symbol originates from the same non-generic component. $\dots 107$
•	Auxiliary unit definitions or local unit definitions may be used to
	avoid repetition of generic unit applications

B.9 Libraries

•	Libraries are named collections of named specifications
•	Local libraries are self-contained
•	Distributed libraries support reuse
•	Different versions of the same library are distinguished by
	hierarchical version numbers
•	Local libraries are self-contained collections of specifications 112
•	Specifications can refer to previous items in the same library 113
•	All kinds of named specifications can be included in libraries
•	Display, parsing, and literal syntax annotations apply to entire
	libraries
•	Libraries and library items can have author and date annotations116
•	Libraries can be installed on the Internet for remote access

•	Validated libraries can be registered for public access
•	Libraries should include appropriate annotations
•	Libraries can include items downloaded from other libraries
•	Substantial libraries of basic datatypes are already available 119
•	Libraries need not be registered for public access
•	Subsequent versions of a library are distinguished by explicit
	version numbers
•	Libraries can refer to specific versions of other libraries
•	All downloadings should be collected at the beginning of a library 122

B.10 Foundations

•	A complete presentation of CASL is in the <i>Reference Manual</i> 125
•	CASL has a definitive summary
•	CASL has a complete formal definition
•	Abstract and concrete syntax of CASL are defined formally126
•	CASL has a complete formal semantics
•	CASL specifications denote classes of models
•	The semantics is largely institution-independent
•	The semantics is the ultimate reference for the meanings of all CASL
	constructs
•	Proof systems for various layers of CASL are provided
•	The foundations of our CASL are rock-solid!

B.11 Tools

•	CASL specifications can be checked for well-formedness using a	
	form-based web page	. 131
•	The Heterogeneous Tool Set (HETS) is the main analysis tool for	
	Casl.	. 132
•	HETS can be used for parsing and checking static well-formedness	
	of specifications.	. 133
•	HETS also displays and manages proof obligations, using	
	development graphs	. 134
•	Nodes in a development graph correspond to CASL specifications.	
	Arrows show how specifications are related by the structuring	
	constructs	. 135
•	Internal nodes in a development graph correspond to unnamed	
	parts of a structured specification	. 137
•	HOL-CASL is an interactive theorem prover for CASL, based on the	
	tactical theorem prover ISABELLE.	. 138
•	CASL is linked to ISABELLE/HOL by an encoding	. 138
•	Asf+SDF was used to prototype the CASL syntax	. 139

210 B Points to Bear in Mind

•	The AsF+SDF Meta-Environment provides syntax-directed editing
	of CASL specifications

B.12 Basic Libraries

•	The Cas	L Basic	Libraries	$\operatorname{contain}$	the	standar	rd o	lataty	pes	 . 143	
	TT	1	1			C 1 D	、 ·	T +1		4 4 0	

• Hets can be used to get an overview of the Basic Libraries. $\ldots \ldots 143$

The Steam-Boiler Control Specification Problem

For completeness, the text describing the steam-boiler control system case study, as originally provided by Jean-Raymond Abrial, is reproduced here (except for the "Additional Information" section, see [1, pp. 507–509]).

C.1 Introduction

This text constitutes an informal specification of a program which serves to control the level of water in a steam-boiler. It is important that the program works correctly because the quantity of water present when the steam-boiler is working has to be neither too low nor to high; otherwise the steam-boiler or the turbine sitting in front of it might be seriously affected.

The proposed specification is derived from an original text that has been written by LtCol. J.C. Bauer for the Institute for Risk Research of the University of Waterloo, Ontario, Canada. The original text has been submitted as a competition problem to be solved by the participants to the International Software Safety Symposium organized by the Institute for Risk Research. It has been given to us by the Institut de Protection et de Sûreté Nucléaire, Fontenay-aux-Roses, France. We would like to thank the author, the Institute for Risk Research, and the Institut de Protection et de Sûreté Nucléaire for their kind permission to use their text.

The text to follow is severely biased to a particular implementation. This is very often the case with industrial specifications that are rarely independent from a certain implementation people have in mind. In that sense, this specification is realistic. Your first formal specification steps could be *much more abstract* if that seems important to you (in particular if your formalism allows you to do so). In other words, you are encouraged to *structure* your specification in a way that is not necessarily the same as the one proposed in what follows. But in any case, you are asked to demonstrate that your specification can be refined to an implementation that is close enough to the functional requirements of the "specification" proposed below.

 $[\]bigodot$ IFIP International Federation for Information Processing 2004

You might also judge that the specification contain some loose ends and inconsistencies. Do not hesitate to point them out and to take yourself some appropriate decisions. The idea, however, is that such inconsistencies should be solely within the *organization* of the system and *not within its physical properties*.

We are aware of the fact that the text to follow does not propose any precise model of the physical evolution of the system, only elementary suggestions. As a consequence, you may have to take some simple, even simplistic, abstract decisions concerning such a physical model.

C.2 Physical Environment

The system comprises the following units:

- the steam-boiler,
- a device to measure the quantity of water in the steam-boiler,
- four pumps to provide the steam-boiler with water,
- four devices to supervise the pumps (one controller for each pump),
- a device to measure the quantity of steam which comes out of the steamboiler,
- an operator desk,
- a message transmission system.

C.2.1 The Steam-Boiler

The steam-boiler is characterized by the following elements:

- A valve for evacuation of water. It serves only to empty the steam-boiler in its initial phase.
- Its total capacity C (indicated in liters).
- The minimal limit quantity M_1 of water (in liters). Below M_1 the steamboiler would be in danger after five seconds, if the steam continued to come out at its maximum quantity without supply of water from the pumps.
- The maximal limit quantity M_2 of waters (in liters). Above M_2 the steamboiler would be in danger after five seconds, if the pumps continued to supply the steam-boiler with water without possibility to evacuate the steam.
- The minimal normal quantity N_1 of water in liters to be maintained in the steam-boiler during regular operation $(M_1 < N_1)$.
- The maximal normal quantity N_2 of water (in liters) to be maintained in the steam-boiler during regular operation $(N_2 < M_2)$.
- The maximum quantity W of steam (in liters/sec) at the exit of the steam-boiler.
- The maximum gradient U_1 of increase of the quantity of steam (in liters/sec/sec).

• The maximum gradient U_2 of decrease of the quantity of steam (in liters/sec/sec).

C.2.2 The Water Level Measurement Device

The device to measure the level of water in the steam-boiler provides the following information:

• the quantity q (in liters) of water in the steam-boiler.

C.2.3 The Pumps

Each pump is characterized by the following elements:

- Its capacity P (in liters/sec).
- Its functioning mode: on or off.
- It's being started: after having been switched on, the pump needs five seconds to start pouring water into the boiler (this is due to the fact that the pump does not balance instantaneously the pressure of the steamboiler).
- It's being stopped: with instantaneous effect.

C.2.4 The Pump Control Devices

Each pump controller provides the following information:

• the water circulates from the pump to the steam-boiler or, on the contrary, it does not circulate.

C.2.5 The Steam Measurement Device

The device to measure the quantity of steam which comes out of the steamboiler provides the following information:

• a quantity of steam v (in liters/sec).

C.2.6 Summary of Constants and Variables

The following table summarizes the various constants or physical variables of the system:

	Unit	Comment				
	Quantity of water in the steam-boiler					
C	liter	Maximal capacity				
M_1	liter	Minimal limit				
M_2	liter	Maximal limit				
N_1	liter	Minimal normal				
N_2	liter	Maximal normal				
	Outcome of steam at the exit of the steam-boiler					
W	liter/sec	Maximal quantity				
U_1	liter/sec/sec	Maximum gradient of increase				
U_2	liter/sec/sec	Maximum gradient of decrease				
	Capacity of each pump					
P	liter/sec	Nominal capacity				
	Current measures					
q	liter	Quantity of water in the steam-boiler				
p	liter/sec	Throughput of the pumps				
v	liter/sec Quantity of steam exiting the steam-boil					

C.3 The Overall Operation of the Program

The program communicates with the physical units through messages which are transmitted over a number of dedicated lines connecting each physical unit with the control unit. In first approximation, the time for transmission can be neglected.

The program follows a cycle and a priori does not terminate. This cycle takes place each five seconds and consists of the following actions:

- Reception of messages coming from the physical units.
- Analysis of informations which have been received.
- Transmission of messages to the physical units.

To simplify matters, and in first approximation, all messages coming from (or going to) the physical units are supposed to be received (emitted) *simultaneously* by the program at each cycle.

C.4 Operation Modes of the Program

The program operates in different modes, namely *initialization*, *normal*, *de-graded*, *rescue*, *emergency stop*.

C.4.1 Initialization Mode

The *initialization* mode is the mode to start with. The program enters a state in which it waits for the message STEAM-BOILER_WAITING to come from the physical units. As soon as this message has been received the program checks whether the quantity of steam coming out of the steam-boiler is really zero. If the unit for detection of the level of steam is defective (that is, when vis not equal to zero), the program enters the *emergency stop* mode. If the quantity of water in the steam-boiler is above N_2 the program activates the value of the steam-boiler in order to empty it. If the quantity of water in the steamboiler is below N_1 then the program activates a pump to fill the steam-boiler. If the program realizes a failure of the water level detection unit it enters the emergency stop mode. As soon as a level of water between N_1 and N_2 has been reached the program send continuously the signal PROGRAM_READY to the physical units until it receives the signal PHYSICAL_UNITS_READY which must necessarily be emitted by the physical units. As soon as this signal has been received, the program enters either the mode *normal* if all the physical units operate correctly or the mode *degraded* if any physical unit is defective. A transmission failure puts the program into the mode *emergency* stop.

C.4.2 Normal Mode

The normal mode is the standard operating mode in which the program tries to maintain the water level in the steam-boiler between N_1 and N_2 with all physical units operating correctly. As soon as the water level is below N_1 or above N_2 the level can be adjusted by the program by switching the pumps on or off. The corresponding decision is taken on the basis of the information which has been received from the physical units. As soon as the program recognizes a failure of the water level measuring unit it goes into rescue mode. Failure of any other physical unit puts the program into degraded mode. If the water level is risking to reach one of the limit values M_1 or M_2 the program enters the mode emergency stop. This risk is evaluated on the basis of a maximal behavior of the physical units. A transmission failure puts the program into emergency stop mode.

C.4.3 Degraded Mode

The *degraded* mode is the mode in which the program tries to maintain a satisfactory water level despite of the presence of failure of some physical unit. It is assumed however that the water level measuring unit in the steamboiler is working correctly. The functionality is the same as in the preceding case. Once all the units which were defective have been repaired, the program comes back to *normal* mode. As soon as the program sees that the water level measuring unit has a failure, the program goes into mode *rescue*. If the

water level is risking to reach one of the limit values M_1 or M_2 the program enters the mode *emergency stop*. A transmission failure puts the program into *emergency stop* mode.

C.4.4 Rescue Mode

The *rescue* mode is the mode in which the program tries to maintain a satisfactory water level despite of the failure of the water level measuring unit. The water level is then estimated by a computation which is done taking into account the maximum dynamics of the quantity of steam coming out of the steam-boiler. For the sake of simplicity, this calculation can suppose that exactly n liters of water, supplied by the pumps, do account for exactly the same amount of boiler contents (no thermal expansion). This calculation can however be done only if the unit which measures the quantity of steam is itself working and if one can rely upon the information which comes from the units for controlling the pumps. As soon as the water measuring unit is repaired, the program returns into mode *degraded* or into mode *normal*. The program goes into *emergency stop* mode if it realizes that one of the following cases holds: the unit which measures the outcome of steam has a failure, or the units which control the pumps have a failure, or the water level risks to reach one of the two limit values. A transmission failure puts the program into *emergency stop* mode.

C.4.5 Emergency Stop Mode

The *emergency stop* mode is the mode into which the program has to go, as we have seen already, when either the vital units have a failure or when the water level risks to reach one of its two limit values. This mode can also be reached after detection of an erroneous transmission between the program and the physical units. This mode can also be set directly from outside. Once the program has reached the *Emergency stop* mode, the physical environment is then responsible to take appropriate actions, and the program stops.

C.5 Messages Sent by the Program

The following messages can be sent by the program:

- MODE(m): The program sends, at each cycle, its current mode of operation to the physical units.
- PROGRAM_READY: In *initialization* mode, as soon as the program assumes to be ready, this message is continuously sent until the message PHYSICAL_UNITS_READY coming from the physical units has been received.

- VALVE: In *initialization* mode this message is sent to the physical units to request opening and then closure of the valve for evacuation of water from the steam-boiler.
- OPEN_PUMP(n): This message is sent to the physical units to activate a pump.
- CLOSE_PUMP(n): This message is sent to the physical units to stop a pump.
- PUMP_FAILURE_DETECTION(n): This message is sent (until receipt of the corresponding acknowledgement) to indicate to the physical units that the program has detected a pump failure.
- PUMP_CONTROL_FAILURE_DETECTION(n): This message is sent (until receipt of the corresponding acknowledgement) to indicate to the physical units that the program has detected a failure of the physical unit which controls a pump.
- LEVEL_FAILURE_DETECTION: This message is sent (until receipt of the corresponding acknowledgement) to indicate to the physical units that the program has detected a failure of the water level measuring unit.
- STEAM_FAILURE_DETECTION: This message is sent (until receipt of the corresponding acknowledgement) to indicate to the physical units that the program has detected a failure of the physical unit which measures the outcome of steam.
- PUMP_REPAIRED_ACKNOWLEDGEMENT(n): This message is sent by the program to acknowledge a message coming from the physical units and indicating that the corresponding pump has been repaired.
- PUMP_CONTROL_REPAIRED_ACKNOWLEDGEMENT(n): This message is sent by the program to acknowledge a message coming from the physical units and indicating that the corresponding physical control unit has been repaired.
- LEVEL_REPAIRED_ACKNOWLEDGEMENT: This message is sent by the program to acknowledge a message coming from the physical units and indicating that the water level measuring unit has been repaired.
- STEAM_REPAIRED_ACKNOWLEDGEMENT: This message is sent by the program to acknowledge a message coming from the physical units and indicating that the unit which measures the outcome of steam has been repaired.

C.6 Messages Received by the Program

The following messages can be received by the program:

- STOP: When the message has been received three times in a row by the program, the program must go into *emergency stop*.
- STEAM_BOILER_WAITING: When this message is received in *initialization* mode it triggers the effective start of the program.

- PHYSICAL_UNITS_READY: This message when received in *initialization* mode acknowledges the message PROGRAM_READY which has been sent previously by the program.
- PUMP_STATE(n, b): This message indicates the state of pump n (open or closed). This message must be present during each transmission.
- PUMP_CONTROL_STATE(n, b): This message gives the information which comes from the control unit of pump n (there is flow of water or there is no flow of water). This message must be present during each transmission.
- LEVEL(v): This message contains the information which comes from the water level measuring unit. This message must be present during each transmission.
- STEAM(v): This message contains the information which comes from the unit which measures the outcome of steam. This message must be present during each transmission.
- PUMP_REPAIRED(n): This message indicates that the corresponding pump has been repaired. It is sent by the physical units until a corresponding acknowledgement message has been sent by the program and received by the physical units.
- PUMP_CONTROL_REPAIRED(n): This message indicates that the corresponding control unit has been repaired. It is sent by the physical units until a corresponding acknowledgement message has been sent by the program and received by the physical units.
- LEVEL_REPAIRED: This message indicates that the water level measuring unit has been repaired. It is sent by the physical units until a corresponding acknowledgement message has been sent by the program and received by the physical units.
- STEAM_REPAIRED: This message indicates that the unit which measures the outcome of steam has been repaired. It is sent by the physical units until a corresponding acknowledgement message has been sent by the program and received by the physical units.
- PUMP_FAILURE_ACKNOWLEDGEMENT(n): By this message the physical units acknowledge the receipt of the corresponding failure detection message which has been emitted previously by the program.
- PUMP_CONTROL_FAILURE_ACKNOWLEDGEMENT(n): By this message the physical units acknowledge the receipt of the corresponding failure detection message which has been emitted previously by the program.
- LEVEL_FAILURE_ACKNOWLEDGEMENT: By this message the physical units acknowledge the receipt of the corresponding failure detection message which has been emitted previously by the program.
- STEAM_FAILURE_ACKNOWLEDGEMENT: By this message the physical units acknowledge the receipt of the corresponding failure detection message which has been emitted previously by the program.
C.7 Detection of Equipment Failures

The following erroneous kinds of behavior are distinguished to decide whether certain physical units have a failure:

- PUMP: (1) Assume that the program has sent a start or stop message to a pump. The program detects that during the following transmission that pump does not indicate its having effectively been started or stopped. (2) The program detects that the pump changes its state spontaneously.
- PUMP_CONTROLLER: (1) Assume that the program has sent a start or stop message to a pump. The program detects that during the second transmission after the start or stop message the pump does not indicate that the water is flowing or is not flowing; this despite of the fact that the program knows from elsewhere that the pump is working correctly. (2) The program detects that the unit changes its state spontaneously.
- WATER_LEVEL_MEASURING_UNIT: (1) The program detects that the unit indicates a value which is out of the valid static limits (that is, between θ and C). (2) The program detects that the unit indicates a value which is incompatible with the dynamics of the system.
- STEAM_LEVEL_MEASURING_UNIT: (1) The program detects that the unit indicates a value which is out of the valid static limits (that is, between θ and W). (2) The program detects that the unit indicates a value which is incompatible with the dynamics of the system.
- TRANSMISSION: (1) The program receives a message whose presence is aberrant. (2) The program does not receive a message whose presence is indispensable.

References

- J.-R. Abrial, E. Börger, and H. Langmaack, editors. Formal Methods for Industrial Applications: Specifying and Programming the Steam Boiler Control, LNCS Vol. 1165. Springer, 1996.
- E. Astesiano, M. Bidoit, B. Krieg-Brückner, P. D. Mosses, D. Sannella, and A. Tarlecki. CASL: The Common Algebraic Specification Language. *Theoretical Comput. Sci.*, 286(2):153–196, 2002.
- E. Astesiano, H.-J. Kreowski, and B. Krieg-Brückner, editors. Algebraic Foundations of Systems Specification. IFIP State-of-the-Art Reports. Springer, 1999.
- 4. S. Autexier, D. Hutter, T. Mossakowski, and A. Schairer. The development graph manager MAYA (system description). In H. Kirchner and C. Ringeissen, editors, Algebraic Methods and Software Technology, 9th International Conference, AMAST 2002, Saint-Gilles-les-Bains, Reunion Island, France, Proceedings, LNCS Vol. 2422, pages 495–502. Springer, 2002.
- S. Autexier and T. Mossakowski. Integrating HOL-CASL into the development graph manager MAYA. In A. Armando, editor, Frontiers of Combining Systems, 4th International Workshop, FroCoS 2002, Santa Margherita Ligure, Italy, Proceedings, LNCS Vol. 2309, pages 2–17. Springer, 2002.
- J. A. Bergstra, J. Heering, and P. Klint. The algebraic specification formalism ASF. In J. A. Bergstra, J. Heering, and P. Klint, editors, *Algebraic Specification*, ACM Press Frontier Series. Addison-Wesley, 1989.
- M. Bidoit. Development of modular specifications by stepwise refinements using the PLUSS specification language. In C. Rattray and R. G. Clark, editors, *Unified Computation Laboratory: Modelling, Specifications, and Tools*, pages 171–192. Oxford Univ. Press, 1992.
- M. Bidoit, C. Chevenier, C. Pellen, and J. Ryckbosch. An algebraic specification of the steam-boiler control system. In Abrial et al. [1], pages 79–108.
- M. Bidoit, M.-C. Gaudel, and A. Mauboussin. How to make algebraic specifications more understandable? An experiment with the PLUSS specification language. *Science of Computer Programming*, 12(1):1–38, 1989.
- M. Bidoit, H.-J. Kreowski, P. Lescanne, F. Orejas, and D. Sannella, editors. Algebraic System Specification and Development. LNCS Vol. 501. Springer, 1991.
- M. Bidoit, D. Sannella, and A. Tarlecki. Architectural specifications in CASL. Formal Aspects of Computing, 13:252–273, 2002.

- M. G. J. van den Brand, H. A. de Jong, P. Klint, and P. Olivier. Efficient annotated terms. Software, Practice & Experience, 30:259–291, 2000.
- M. G. J. van den Brand, A. van Deursen, J. Heering, H. A. de Jong, M. de Jonge, T. Kuipers, P. Klint, L. Moonen, P. A. Olivier, J. Scheerder, J. J. Vinju, E. Visser, and J. Visser. The AsF+SDF Meta-Environment: A componentbased language development environment. In R. Wilhelm, editor, *Compiler Construction, 10th International Conference, CC 2001, Genova, Italy, Proceedings*, LNCS Vol. 2027, pages 365–370. Springer, 2001.
- M. Broy, C. Facchi, R. Grosu, R. Hettler, H. Hußmann, D. Nazareth, F. Regensburger, O. Slotosch, and K. Stølen. The requirement and design specification language SPECTRUM: An informal introduction (v 1.0). Technical Report TUM-I9311, TUM-I9312, Institut für Informatik, Technische Universität München, 1993.
- R. M. Burstall and J. A. Goguen. The semantics of CLEAR, a specification language. In D. Bjørner, editor, Abstract Software Specifications, 1979 Copenhagen Winter School, Proceedings, LNCS Vol. 86, pages 292–332. Springer, 1980.
- M. Cerioli, M. Gogolla, H. Kirchner, B. Krieg-Brückner, Z. Qian, and M. Wolf, editors. Algebraic System Specification and Development: Survey and Annotated Bibliography. BISS Monographs. Shaker, 2nd edition, 1997.
- M. Cerioli and G. Reggio, editors. Recent Trends in Algebraic Development Techniques, 15th International Workshop, WADT 2001, Joint with the CoFI WG Meeting, Genova, Italy, 2001, Selected Papers, LNCS Vol. 2267. Springer, 2001.
- F. Chen, G. Rosu, and R. P. Venkatesan. Rule-based analysis of dimensional safety. In R. Nieuwenhuis, editor, *Rewriting Techniques and Applications*, 14th International Conference, RTA 2003, Valencia, Spain, Proceedings, LNCS Vol. 2706, pages 197–207. Springer, 2003.
- I. Claßen, H. Ehrig, and D. Wolz. Algebraic Specification Techniques and Tools for Software Development. AMAST Series in Computing Vol. 1. World Scientific, 1993.
- 20. CoFI (The Common Framework Initiative). CASL *Reference Manual*. LNCS, IFIP Series. Springer, 2004. To appear.
- 21. CoFI (The Common Framework Initiative) Tools Group. Home page. http://www.cofi.info/Tools.
- A. van Deursen, J. Heering, and P. Klint, editors. Language Prototyping: An Algebraic Specification Approach. AMAST Series in Computing Vol. 5. World Scientific, 1996.
- J. A. Goguen and R. M. Burstall. Institutions: Abstract model theory for specification and programming. J. ACM, 39(1):95–146, 1992.
- 24. J. A. Goguen, T. Winkler, J. Meseguer, K. Futatsugi, and J.-P. Jouannaud. Introducing OBJ3. In J. Goguen and G. Malcolm, editors, *Software Engineering with OBJ: Algebraic Specification in Action.* Kluwer, 1992.
- J. V. Guttag and J. J. Horning. Larch: Languages and Tools for Formal Specification. Springer, 1993.
- S. Kahrs, D. Sannella, and A. Tarlecki. The definition of Extended ML: A gentle introduction. *Theoretical Comput. Sci.*, 173:445–484, 1997.
- J. Loeckx, H.-D. Ehrich, and M. Wolf. Specification of Abstract Data Types. Wiley/Teubner, 1996.
- 28. T. Mossakowski. CASL: From semantics to tools. In S. Graf and M. Schwartzbach, editors, *Tools and Algorithms for the Construction and Analysis of Sys*-

tems, 6th International Conference, TACAS 2000, Berlin, Germany, Proceedings, LNCS Vol. 1785, pages 93–108. Springer, 2000.

- 29. T. Mossakowski. Relating CASL with other specification languages: The institution level. *Theoretical Comput. Sci.*, 286:367–475, 2002.
- T. Mossakowski, Kolyang, and B. Krieg-Brückner. Static semantic analysis and theorem proving for CASL. In F. Parisi-Presicce, editor, *Recent Trends in Algebraic Development Techniques*, 12th International Workshop, WADT'97, Tarquinia, Italy, 1997, Selected Papers, LNCS Vol. 1376, pages 333–348. Springer, 1998.
- L. C. Paulson. Isabelle: A Generic Theorem Prover. LNCS Vol. 828. Springer, 1994.
- M. Roggenbach and L. Schröder. Towards trustworthy specifications I: Consistency checks. In Cerioli and Reggio [17], pages 305–327.
- D. Sannella. The Common Framework Initiative for algebraic specification and development of software: Recent progress. In Cerioli and Reggio [17], pages 328–343.
- 34. D. Sannella and A. Tarlecki. Foundations of Algebraic Specification and Formal Program Development. To appear.
- 35. D. Sannella and A. Tarlecki. Essential concepts of algebraic specification and program development. *Formal Aspects of Computing*, 9:229–269, 1997.
- M. Wirsing. Structured algebraic specifications: A kernel language. *Theoretical Comput. Sci.*, 42:123–249, 1986.
- M. Wirsing. Algebraic specification. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 13. Elsevier Science Publishers, Amsterdam; and MIT Press, 1990.

List of Named Specifications

STRICT_ PARTIAL_ ORDER	24
Total_ Order	25
TOTAL_ORDER_ WITH_ MINMAX	26
VARIANT_ OF_ TOTAL_ ORDER_ WITH_ MINMAX	27
Partial_ Order	27
Partial_Order_1	28
Implies_ Does_ Not_ Hold	28
Monoid	29
Generic_ Monoid	29
Non_ Generic_ Monoid	30
GENERIC_ COMMUTATIVE_ MONOID	30
GENERIC_ COMMUTATIVE_ MONOID_ 1	31
Container	32
Marking_ Container	32
Generated_ Container	33
Generated_ Container_ Merge	34
Generated_ Set	35
NATURAL	36
Color	37
INTEGER	37
NATURAL_ORDER	38
NATURAL_ARITHMETIC	38
INTEGER_ ARITHMETIC	39
INTEGER_ ARITHMETIC_ ORDER	40
LIST	40
Set	40
TRANSITIVE_ CLOSURE	41
NATURAL_ WITH_ BOUND	41
Set_ Choose	42
Set_ Generated	42
Set_ Union	43

Set_ Union_ 1	44
UNNATURAL	44
Set_ Partial_ Choose	47
Set_ Partial_ Choose_ 1	50
Set_ Partial_ Choose_ 2	50
NATURAL WITH BOUND AND ADDITION	51
Set_ Partial_ Choose_ 3	51
NATURAL_PARTIAL_PRE	52
NATURAL_ PARTIAL_ SUBTRACTION_ 1	52
NATURAL_ PARTIAL_ SUBTRACTION.	52
LIST SELECTORS 1	53
LIST SELECTORS 2	53
LIST SELECTORS	54
NATURAL SUC PRE	54
PAIR 1	54
PART CONTAINER	54
NATURAL PARTIAL SUBTRACTION 2	55
CENERIC MONOID 1	57
VEHICI F	58
MODE VEHICLE	50
SDEED RECULATION	50
NATURAL SUBSORTS	60
DOSTINE	61
	61
POSITIVE PRE	62
NATURAL DOCUMER ADJUNETIC	62
INTEGED ADJUNETIC 1	64
SET EDDOD CHOOSE	65
SEL ERROR CHOOSE	65
JEI_ERROR_CHOOSE_I	67
LIST_DEL	69
CDT TO LIT	60
SET_ TO_ LIST	00
STACK	09
LIST_SET_1	(1
NATURAL PARTIAL SUBTRACTION 3	(1
NATURAL_PARTIAL_SUBTRACTION_4	(1
PARTIAL_ORDER_ 2	72
LIST_ORDER	73
LIST_ URDER_ SORTED	73
WRONG_LIST_ORDER_SORTED	74
LIST_ ORDER_ SORTED_ 2	74
LIST_ URDER_ SORTED_ 3	75
LIST_ URDER_ NAT	78
NAT_ WORD	79
NAT_ WORD_ 1	79

This_ Is_ Wrong	80
LIST_ ORDER_ POSITIVE	81
NAT_ WORD_ 2	81
PAIR	82
Homogeneous_ Pair_ 1	82
Homogeneous_ Pair	82
TABLE	82
PAIR_ NATURAL_ COLOR	82
PAIR_ NATURAL_ COLOR_ 1	83
PAIR_ NATURAL_ COLOR_ 2	83
PAIR_ Pos	83
Pair_ Pos_ 1	83
My_ Table	83
Set_ of_ List	84
MISTAKE	84
Set_ and_ List	84
This_ Is_ Still_ Wrong	84
List_ Rev	85
List_ Rev_ Nat	85
Two_ Lists	85
Two_ Lists_ 1	86
Monoid_ C	86
Monoid_ of_ Monoid	86
List_ Rev_ Order	87
List_ Rev_ with_ Two_ Orders	87
List_Weighted_Elem	88
List_Weighted_ Pair_ Natural_ Color	89
List_ Weighted_ Instantiated	89
List_ Length	89
List_ Length_ Natural	90
Integer_ as_ Total_ Order	90
Integer_ as_ Reverse_ Total_ Order	91
List_Rev_with_Two_Orders_1	91
List_ as_ Monoid	91
Елем	94
Cont	94
Cont_ Diff	95
Req	95
FLAT_ REQ	95
System	96
System_ 1	97
Cont_ Diff_ 1	99
Inconsistent	99
System_ G	100
System_ A	101

CONT_ COMP	102
DIFF_ COMP	102
System_ G1	102
DIFF_ COMP_ 1	102
Other_ System	103
Other_ System_ 1	104
Set_ Comp	105
Cont2Set	105
Arch_ Cont2Set_ Nat	105
Arch_ Cont2Set	106
Arch_ Cont2Set_ Used	106
Arch_ Cont2Set_ Nat_ 1	107
WRONG_ ARCH_ SPEC	108
BADLY_STRUCTURED_ARCH_SPEC	108
Well_Structured_Arch_Spec	109
ANOTHER_ WELL_ STRUCTURED_ ARCH_ SPEC	109
NATURAL_ORDER_ 2	134
v1	134
v2	137
NUMBERS	144
BASIC/NUMBERS	145
NAT	145
Імт	146
RAT	.147
STRUCTURED DATATYPES	147
BASIC/STRUCTURED DATATYPES.	147
GENERATESET	148
SET	148
GenerateMap	148
Мар	149
FINITE	149
ΤοταιΜαρ	149
GENERATEBAG	150
BAG	150
GENERATELIST	150
LIST	150
Array	151
GENERATERINTREE	152
BINTREE	152
GENERATERINTREE?	152
RINTDEF9	152
CENEDATEKTDEE	152
KTRFF	152
Cened ateNTdee	15/
NTDEF	15/
1 X ± 101212 · · · · · · · · · · · · · · · · · ·	- TO4

Value
Basics
Messages_ Sent
Messages_ Received
SBCS_ CONSTANTS
Preliminary
SBCS
SBCS_ STATE_ 1
Mode_ Evolution
SBCS_ STATE_ 2
STATUS_ EVOLUTION
Message_ Transmission_ System_ Failure
SBCS_ STATE_ 3
Pump_ Failure
SBCS_ STATE_ 4
Pump_ Controller_ Failure174
SBCS_ STATE_ 5
Steam_ Failure
Level_ Failure
Failure_ Detection
Steam_ And_ Level_ Prediction
PUMP_ STATE_ PREDICTION180
PUMP_ CONTROLLER_ STATE_ PREDICTION
PU_ PREDICTION
SBCS_ ANALYSIS
SBCS_ STATE
Steam_ Boiler_ Control_ System
Arch_ Sbcs
Arch_ Preliminary187
SBCS_ STATE_ IMPL
UNIT_ SBCS_ STATE
Arch_ Analysis
Arch_ Failure_ Detection
Arch_ Prediction

Index of Library and Specification Names

ANOTHER_WELL_STRUCTURED_ARCH_ Spec 109 Arch_Analysis 188 Arch_Cont2Set 106 Arch_Cont2Set_Nat 105 Arch_Cont2Set_Nat_1 107 Arch_Cont2Set_Used 106 ARCH_FAILURE_DETECTION 189Arch_Prediction 189 Arch_Preliminary 187 Arch_Sbcs 186 Array 151 BADLY_STRUCTURED_ARCH_SPEC 108 BAG 150 BASIC/NUMBERS 145 **BASIC/STRUCTUREDDATATYPES** 147BASICS 159 BINTREE 152 BINTREE2 152 Color 37 Cont 94 Cont2Set 105 CONT_COMP 102 CONT_DIFF 95 CONT_DIFF_1 99 Container 32 DIFF_COMP 102 DIFF_COMP_1 102 ELEM 94

FAILURE_DETECTION 175 FINITE 149 FLAT_REQ 95 GenerateBag 150 GENERATEBINTREE 152 GENERATEBINTREE2 152 GENERATED_CONTAINER 33 GENERATED_CONTAINER_MERGE 34 Generated_Set 35 GenerateKTree 153 GenerateList 150 GENERATEMAP 148 GENERATENTREE 154 GenerateSet 148 GENERIC_COMMUTATIVE_MONOID 30 GENERIC_COMMUTATIVE_MONOID_1 31Generic_Monoid 29 Generic_Monoid_1 57 Homogeneous_Pair 82 Homogeneous_Pair_1 82 IMPLIES_DOES_NOT_HOLD 28 Inconsistent 99 INT 146 INTEGER 37 INTEGER_ARITHMETIC 39 INTEGER_ARITHMETIC_1 64 INTEGER_ARITHMETIC_ORDER 40 INTEGER_AS_REVERSE_TOTAL_ORDER 91 INTEGER_AS_TOTAL_ORDER 90

KTREE 153

LEVEL_FAILURE 175 LIST 40, 150 LIST_AS_MONOID 91 LIST_CHOOSE 68 LIST_LENGTH 89 LIST_LENGTH_NATURAL 90 LIST_ORDER 73 LIST_ORDER_NAT 78 LIST_ORDER_POSITIVE 81 LIST_ORDER_SORTED 73 LIST_ORDER_SORTED_2 74 LIST_ORDER_SORTED_3 75 LIST_REV 85 LIST_REV_NAT 85 LIST_REV_ORDER 87 LIST_REV_WITH_TWO_ORDERS 87 LIST REV WITH TWO ORDERS 1 91 LIST_SELECTORS 54 LIST_SELECTORS_1 53 LIST_SELECTORS_2 53 LIST_SET 67 LIST_SET_1 71 LIST_WEIGHTED_ELEM 88 LIST_WEIGHTED_INSTANTIATED 89 LIST_WEIGHTED_PAIR_NATURAL_ Color 89 MAP 149 Marking_Container 32 Message_Transmission_System_ FAILURE 171 Messages_Received 159 Messages_Sent 159 MISTAKE 84 MODE_EVOLUTION 165 Monoid 29 Monoid_C 86 Monoid_of_Monoid 86 More_Vehicle 59 My_TABLE 83 Nat 145 NAT_WORD 79 NAT_WORD_1 79 NAT_WORD_2 81 NATURAL 36

NATURAL_ARITHMETIC

38

NATURAL_ORDER 38 NATURAL_ORDER_2 134 NATURAL_PARTIAL_PRE 52 NATURAL_PARTIAL_SUBTRACTION 52 NATURAL_PARTIAL_SUBTRACTION_1 52 NATURAL PARTIAL SUBTRACTION 2 55NATURAL_PARTIAL_SUBTRACTION_3 71NATURAL_PARTIAL_SUBTRACTION_4 71NATURAL_POSITIVE_ARITHMETIC 62 NATURAL_SUBSORTS 60 NATURAL_SUC_PRE 54 NATURAL_WITH_BOUND 41 NATURAL_WITH_BOUND_AND_ Addition 51 Non_Generic_Monoid 30 NTREE 154 NUMBERS 144 OTHER_SYSTEM 103 OTHER_SYSTEM_1 104 PAIR 82 PAIR_1 54 PAIR_NATURAL_COLOR 82 PAIR_NATURAL_COLOR_1 83 PAIR_NATURAL_COLOR_2 83 PAIR_POS 83 PAIR_POS_1 83 Part_Container 54 PARTIAL_ORDER 27 PARTIAL_ORDER_1 28 PARTIAL_ORDER_2 72 Positive 61 POSITIVE_ARITHMETIC 61 POSITIVE_PRE 62 Preliminary 160 PU_PREDICTION 181 PUMP_CONTROLLER_FAILURE 174PUMP_CONTROLLER_STATE_ PREDICTION 181 PUMP_FAILURE 173 PUMP_STATE_PREDICTION 180 Rat 147 Req 95

SBCS 161 SBCS_ANALYSIS 182 SBCS_CONSTANTS 160 SBCS_STATE 183 SBCS_STATE_1 163SBCS_STATE_2 169SBCS_STATE_3 172 SBCS_STATE_4 173SBCS_STATE_5 174 SBCS_STATE_IMPL 187 Set 40, 148 Set_and_List 84 Set_Choose 42 Set_Comp 105 Set_Error_Choose 65 Set_Error_Choose_1 65Set_Generated 42 Set_of_List 84 Set_Partial_Choose 47 Set_Partial_Choose_1 50Set_Partial_Choose_2 50Set_Partial_Choose_3 51Set_to_List 68 Set_Union 43 Set_Union_1 44 Speed_Regulation 59Stack 69 STATUS_EVOLUTION 169 STEAM_AND_LEVEL_PREDICTION 178 STEAM_BOILER_CONTROL_SYSTEM 183

STEAM_FAILURE 175 Strict_Partial_Order 24StructuredDatatypes 147 System 96 System_1 97 System_A 101 System_G 100 System_G1 102 TABLE 82 THIS_IS_STILL_WRONG 84 THIS_IS_WRONG 80 TOTAL_ORDER 25 TOTAL_ORDER_WITH_MINMAX 26 TotalMap 149 TRANSITIVE_CLOSURE 41 Two_Lists 85 Two_Lists_1 86 UNIT_SBCS_STATE 188 UNNATURAL 44 v1 134 v2 137 VALUE 158 VARIANT_OF_TOTAL_ORDER_WITH_ MinMax 27 Vehicle 58 Well_Structured_Arch_Spec 109 WRONG_ARCH_SPEC 108 WRONG_LIST_ORDER_SORTED 74

Concept Index

abbreviation 23 abstract syntax 126 algebraic signature 12analysis, static 127, 133 annotation 115 associativity 115author 116 date 116.121 definitional extension 43display 27, 115, 119 implies 28 label 25 literal syntax 118 parsing 115, 119 precedence 115relative precedence 115 architectural specification 93, 186 argument specification 19, 30fitting 77 ASF+SDF 139assertion definedness 50 subsort membership 59 associativity annotation 115 attribute 29,116 ATerms 137, 139 attribute 29 associativity 29,116 commutativity 2929idempotence unit 29

author annotation 116 auxiliary operation 71 predicate 71symbol 73 axiom 11, 12, 14, 25 basic datatypes, libraries of 117, 119, 143 specification 11, 194 body 30 specification 18carrier set 12 casting 64 Cats 138 class of models 12 closed specification 114 system 106 world assumption 38 comment 25 commutativity attribute 29compatibility 105, 106 overloading with embedding 62 renaming 70 subsorting 14 completeness 128component -93declaration 96 generic 100 application of 101 specification 96

named 102 composition 19 generic specification 84 compound sort 12,85 symbol 85 operation 87 predicate 87 concrete syntax 126 consequence relation 128conservative extension 19,98 consistency checker 141 consistent 12 constant operation 13 overloaded 31 symbol 29 constraint 11,12 sort generation 15, 33 constructor 32 partial 55 current version 121 datatype declaration 32.59enumerated 37 free 37 generated 33 structured 147 date annotation 116, 121 decimal notation 118 declaration 11 before use *see* linear visibility component 96 datatype 32, 59 function partial 47 total 26 subsort 57, 59 symbol 11 unit 96 variable 25 global 27 decomposition 19 deduction rule 128 definedness assertion 50definition domain of 13.50operation 26

predicate 28 style 44 subsort 60 unit 109 definitional extension 43 development graph 134 disambiguation 31 display annotation 27, 115, 119 format 27 distributed library 111, 116 domain of definition 13, 50semantic 127 downloaded specification 112, 118 editor, syntax-directed 139 embedding 13,64 compatibility with overloading 62 end-of-line comment 25 enumerated datatype 37 equation existential 15,55 strong 15, 48, 55 error supersort 65 evolution of specification 112existential equation 15,55 quantification 14,25 expansion 16 explicit fitting 81 exported symbol 16,25 expression, unit 106 extension 25conservative 19,98 definitional 43 specification 18,68 fitting argument specification 77 explicit 81 implicit 80 morphism 19 symbol map 81 fixed part 88 formal specification 156 verification 128, 138 formalization process 156

formula 25 free datatype 37 model 18 specification 18,36 freeness constraint 18, 37 function partial 13,47 persistent 19,100total 13, 23 generated datatype 33 specification 33 generic component 100 application 101 specification 15, 18, 30, 77 composition 84 view 91 global variable declaration 27Hets 132hiding specification 17,71 with revealing 72 HOL-CASL 138 homomorphism 16 idempotence attribute 29identity renaming 71 ill-formed instantiation 80,90 implementation 93, 140 implicit fitting 80 implies annotation 28imports 19,88 inconsistent 12 independence, institution 127, 137induction 34 proof 139 scheme 35 inductive definition 38 style 44 informal requirements 156 inherited 58 initial model 16.38input format 27instantiation 15, 19, 30, 77, 78 ill-formed 80,90

institution 12, 127, 128 independence 127, 137 interpretation 11, 12 isomorphism 16 junk 15 label annotation 25 libraries of basic datatypes 117, 119, 143 library 111 current version 121 distributed 111, 116 item 114 local 111, 112 registered 117 self-contained 111, 112 under development 120 validation 117 version 112, 121linear visibility 26, 96, 113 literal syntax annotation 118 local library 111, 112 specification 73unit definition 109 logic 15, see institution loose specification 24 many-sorted signature 12 map, symbol 70 mathematical theory 126 Maya 138meaning 12 mixfix notation 24 model class 12 free 18 of specification 12 semantics 127 morphism fitting 19 signature 16 specification 16 multi-line comment 25mutual recursion 36, 113 name, path 117 named specification 15, 24, 75

view 19 natural semantics 127 negation as failure 38 'no junk, no confusion' principle 37 'no junk' principle 33 non-isomorphic models 25non-linear visibility 36,63 notation, decimal 118 number 144 observer 43 definition style 44 open system 106 operation auxiliary 71 boolean-valued v. predicate 13, 18 constant 13 definition 26 symbol 13 compound 87 overloaded constant 31 symbol 31 renaming 70 overloading 14,61 compatibility with embedding 62 with renaming 70 parameter 30 specification 18, 77, 78 parsing 133, 137, 139 annotation 115, 119 partial constructor 55 function 13,47 declaration 47 selector 54 path name 117 persistent function 19, 100 place-holder 24 precedence annotation 115 predicate auxiliary 71 definition 28 symbol 13, 24 compound 87 v. boolean-valued operation 13, 18 principle

'no junk, no confusion' 37 'no junk' 33 'same name, same thing' 17, 31, 113 process of formalization 156profile 13programming language 140 proof calculus 128 induction 139obligation 28,134 propagate undefinedness 48 quantification 14, 25existential unique-existential 14 universal 14,25 recursion, mutual 36, 113 reduct 16 reference to name 15refinement 156 registered library 117 relative precedence annotation 115 renaming compatibility with overloading 70 identity 71 overloaded symbol 70specification 69 requirements informal 156 specification 156 result sort 13 reuse of specification 111 revealing with hiding 72rewriting 140 'same name, same thing' principle 17. 31, 113 satisfaction relation 12 scheme, induction 35 selector 53 partial 54 total 54 self-contained library 111, 112 specification *see* closed semantic domain 127semantics 12, 126 model 127 natural 127

static 127 set carrier 12 inclusion 13 signature 12 algebraic 12 many-sorted 12 morphism 16 subsorted 12 software modules 94 sort 12,24 compound 12,85generation constraint 15.33result 13 union 59 soundness 128 specification architectural 93, 186 argument 19,30 fitting 77 basic 11.194body 18 closed 114 downloaded 112.118 evolution 112 extension 18,68 fitting argument 77 formal 156 free 18, 36 generated 33 generic 15, 18, 30, 77 hiding 17,71 instantiation 77,78 local -73 loose 24 model of 12 morphism 16 named 15, 24, 75 of component 96 named 102 parameter 18, 77, 78 reference 15 renaming 69 requirements 156 reuse 111 structured 67 subsorted 57 translation 17 union 17,68

unit -98 static analysis 127, 133 semantics 127 strong equation 15, 48, 55 structured datatype 147 specification 67 subsort 13, 57 declaration 57, 59 definition 60 membership assertion 59subsorted signature 12 specification 57 supersort 13, 57 symbol 12 auxiliary 73 compound 85 constant 29 declaration 11 exported 16,25 map 70 fitting 81 operation 13 overloaded 31 predicate 13,24 syntax 126 -directed editor 139 system 106 term, unit 100 theorem prover 138total function 13, 23 declaration 26 selector 54 translation 17 two-valued logic 48 undefined value 48 union sort 59 specification 17,68 unique-existential quantification 14 unit attribute 29 declaration 96 definition 109

expression 106 local definition 109 specification 98 term 100 universal quantification 14, 25 validation 156 of library 117 variable 14 declaration 25 version

```
control 112, 120
current 121
library 121
number 121
of library 112
view 90
generic 91
named 19
visibility
component name 97
linear 26, 96, 113
```