CoreASM: AN EXTENSIBLE MODELING FRAMEWORK & TOOL ENVIRONMENT FOR HIGH-LEVEL DESIGN AND ANALYSIS OF DISTRIBUTED SYSTEMS

by

Roozbeh Farahbod

B.Sc., Sharif University of Technology, 2001

M.Sc., Simon Fraser University, 2004

A Thesis submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
in the School
of
Computing Science

SIMON FRASER UNIVERSITY
Summer 2009

APPROVAL

Name: Roozbeh Farahbod

Degree: Doctor of Philosophy

Title of Thesis: CoreASM: An Extensible Modeling Framework & Tool Environment for High-level Design and Analysis of Distributed Systems

Examining Committee:

Dr. Dirk Beyer, Assistant Professor
Chair

Dr. Uwe Glässer, Professor
Senior Supervisor

Dr. Robert D. Cameron, Professor
Supervisor

Dr. Lou Hafer, Professor
SFU Examiner

Dr. Egon Börger, External Examiner,
Professor of Computer Science,
University of Pisa, Italy

Date Approved: May 12 2009

Declaration of
Partial Copyright Licence

Abstract

Model-based systems engineering naturally requires abstract executable specifications to facilitate simulation and testing in early stages of the system design process. Abstraction and formalization provide effective instruments for establishing critical system requirements by precisely modeling the system prior to construction so that one can analyze and reason about specification and design choices and better understand their implications. There are many approaches to formal modeling of software and hardware systems. Abstract State Machines, or ASMs, are well known for their versatility in computational and mathematical modeling of complex distributed systems with an orientation toward practical applications. They offer a good compromise between declarative, functional and operational views towards modeling of systems. The emphasis on *freedom of abstraction* in ASMs leads to intuitive yet accurate descriptions of the dynamic properties of systems. Since ASMs are in principle executable, the resulting models are validatable and possibly falsifiable by experiment. Finally, the well-defined notion of *step-wise refinement* in ASMs bridges the gap between abstract models and their final implementations.

There is a variety of tools and executable languages available for ASMs, each coming with their own strengths and limitations. Building on these experiences, this work puts forward the design and development of an extensible and executable ASM language and tool architecture, called CoreASM, emphasizing *freedom of experimentation* and *design exploration* in the early phases of the software development process. CoreASM aims at preserving the very idea of ASM modeling—the design of accurate abstract models at the level of abstraction determined by the application domain, while encouraging rapid prototyping of such abstract models for testing and design space exploration. In addition, the extensible language and tool architecture of CoreASM facilitates integration of domain specific concepts and special-purpose tools into its language and modeling environment.

CoreASM has been applied in a broad scope of R&D projects, spanning maritime surveillance, situation analysis, and computational criminology. In light of these applications, we argue that the design and implementation of CoreASM accomplishes its goals; it not only preserves the desirable characteristics of abstract mathematical models, such as conciseness, simplicity and intelligibility, but it also adheres to the methodological guidelines and best practices for ASM modeling.

Keywords: CoreASM; Abstract State Machines; Specification Languages; Executable Specification; Distributed Systems; High-level Design

To Maryam and Marjan,
my darling sisters.

"He who would learn to fly one day, must first learn to stand and walk and run and climb and dance; one cannot fly into flying."

— Friedrich W. Nietzsche

Acknowledgments

I am mostly grateful to my senior supervisor Dr. Uwe Glässer for his enthusiasm, friendship, and generous support and supervision throughout this work. He offered the original idea of CoreASM and provided the motivation for this project to take shape.

I would like to express my sincere gratitude to my dear friend and mentor, Dr. Vincenzo Gervasi. The success of this work would not have been possible without his kind support, encouragement, and inspiring ideas.

I would like to thank Dr. Robert Cameron, Dr. Lou Hafer, and Dr. Tom Shermer for their valuable feedback and inspiring discussions that led to the improvement of this thesis. I would also like to specially thank Dr. Egon Börger for his thorough examination of this thesis and his suggestions, corrections, and remarks on the theoretical and practical aspects of this work.

Many ideas in this work are the outcome of lengthy discussions (and at times heated arguments) with my dear friends and colleagues Mashaal Memon and George Ma. I am grateful for having the opportunity of meeting them in the course of this project.

I would like to express my gratitude to Michael Altenhofen at SAP Research in Germany. During the final phases of this work, Michael offered many suggestions for improvements and invaluable feedback on the practical issues and usability of CoreASM.

My heartfelt thanks goes to my many friends, colleagues, and my family (past and present) who made these years a pleasant and memorable period in my life.

I would also like to acknowledge the people in the School of Computing Science, the administrative and the technical support staff, and the Network Support Group for making this school a productive environment for research and studies. Finally, I also wish to thank the Natural Sciences and Engineering Research Council of Canada (NSERC) and Precarn's Intelligent Systems program for their financial support in the course of this project.

Contents

List of Tables

List of Figures

List of Programs

Part I

Introduction

Chapter 1

Background and Motivation

Computer-based systems are increasingly integrated into our day-to-day life. They either control or provide platforms for our communication networks, transportation facilities, economic markets, health-care systems, and safety and security facilities. With the increasing complexity of these systems, efficient design and development of high quality computational systems that faithfully conform to their requirements are extremely challenging and the costs of design flaws and system failures are high. Proper understanding of the requirements, precisely documenting design decisions, and effectively communicating such decisions with the domain experts as early as possible play important roles in the design of complex systems. These challenges call for adoption of proper engineering methods and tools and have motivated the use of *formal methods* in software engineering.

Abstraction and formalization provide effective instruments for establishing critical system requirements by precisely modeling systems prior to construction so that one can analyze and reason about specification and design choices and better understand their implications [9]. There are many approaches to formal modelling of software and hardware systems. *Abstract State Machines (ASMs)* [25] are well known for their versatility in computational and mathematical modelling of complex distributed systems with an orientation toward practical applications. The ASM framework offers a universal model of computation and serves as an effective instrument for analyzing and reasoning about complex semantic properties of discrete dynamic systems. For almost two decades, abstract state machines have been studied, practiced, and applied in modeling and specification of systems to bridge the gap between formal and pragmatic approaches. Combining common abstraction principles from computational logic, discrete mathematics, and the concept of transition systems,

ASMs have become a well-known method and assumed a major role in providing a solid and flexible mathematical framework for specification and modeling of virtually all kinds of discrete dynamic systems.

In addition, machine assistance plays an increasingly important role in making design and development of complex systems feasible. Abstract executable specifications serve as a basis for design exploration and experimental validation through simulation and testing. Model checking tools based on formal verification techniques help with proving critical properties of systems and assuring "correctness" before deployment.

There is a variety of tools and executable languages available for ASMs, each coming with their own strengths and limitations. In this work, we critically look into their interesting features and potential shortcomings with the goal of understanding the requirements of a modeling language and tool environment that would support high-level design and experimental validation of abstract machine models at the early stages of design and development. Building on these experiences, this work puts forward the design and development of an extensible and executable ASM language and tool architecture, called CoreASM, emphasizing *freedom of experimentation* and *design exploration* in the early phases of the software development process. CoreASM aims at preserving the very idea of ASM modeling—the design of accurate abstract models (*ground models* [17]) at the level of abstraction determined by the application domain, while encouraging rapid prototyping of such models for conformance testing, design space exploration, and experimental validation.

The rest of this introductory chapter is organized as follows. Section 1.1 briefly looks into the concept of a *model* and explores various formal languages and techniques for modeling software intensive computer-based systems. Section 1.2 outlines different approaches toward establishing formal semantics of languages—i.e., what makes modeling languages "formal"—and argues why such formal semantics are necessary. The objectives of this thesis and a brief discussion of the proposed solution are sketched in sections 1.3 and 1.4. Finally, this chapter ends with an outline of the thesis in Section 1.5.

1.1 Modeling Languages

A *model* of a system is an abstract representation of that system so that one can view, manipulate, and reason about it [98]. Such a representation also helps in understanding the complexity that is inherent in the system under study. We build models to increase

productivity, since it is often cheaper to explore and to manipulate the model than the real system. A "good" model omits unnecessary information but accurately reflects the essential aspects of the subject matter in order to help the viewer to clearly see the subject and focus on those essential aspects. A model that is easily understandable can also serve as a means of communication by clearly illustrating the subject and its main concepts and ideas.

There are many modeling languages available to express computational models, each one focusing on certain aspects or targetting certain types of systems. A popular example of a widely used modeling language is the *Unified Modeling Language*, or UML[1] for short. UML is a visual (graphical) language and is one of the most common industrial modeling languages in the area of software engineering. However, UML is an informal language[2] as its semantics is not formally defined.[3]

Our work, however, is focused on the practical application of *formal methods*. According to Daniel M. Berry [9], a formal method is any attempt to use mathematics in the development of a software intensive computer-based system in order to improve the quality of the resulting system. We define a *formal modeling language* as a modeling language that has a formally defined (read "mathematically defined") syntax and a formally defined semantics for that syntax. There are many formal languages and notations for modeling and specification[4] of systems, such as: the *Vienna Development Method (VDM)* [11], one of the longest-established formal methods for the functional modeling of computer-based systems; the family of *Algebraic Specification* languages, currently subsumed under the *Common Algebraic Specification Language*, or *CASL* [10], which are all based on first-order logic with induction, viewing states of systems as first-order many-sorted structures; the family of *process calculi* or *process algebras* languages and approaches to formal modeling of concurrent systems (such as *π-calculus* [105] or *Communicating Sequential Processes (CSP)* [79]), supporting high-level description of interactions between a collection of independent processes; *Specification and Description Language (SDL)* [42], a standard formal language [83]

[1] http://www.uml.org

[2] Some people claim that UML is a semi-formal modeling language since its (abstract) syntax is precisely defined [65].

[3] There have been attempts to formally define the semantics of UML (see [65] for example).

[4] There is a slight difference between a 'model' and a 'specification' of a system. Strictly speaking, a specification of a system tends to view the system as a black box, focusing on the behavior of the system as a whole and its interface to its environment [42, 82]; i.e., focusing on *what* the system does. A model of a system, on the other hand, can include both a specification and a description of the system; i.e., describing *what* the system does and *how* it does it.

for specification and description of reactive and distributed systems, which provides both graphical and textual representations; the *Petri nets*[5] [111] graphical language for description of distributed systems in form of bipartite graphs[6]; the *B method* [2], an abstract machine modeling approach mostly used in the development of software with a rich set of commercially available tools for specification, design, proof and code generation; the *Z notation* [116], a formal specification language for modeling computing systems and formulation of proofs about the intended program behavior based on axiomatic set theory, lambda calculus, and first-order predicate logic; the *Alloy* specification language [84], a light-weight formal specification language (inspired by the Z notation) together with a tool designed for providing fully automatic analysis of software specifications; and last but not least, the ASM method [25], a versatile semantic framework for computational modeling of virtually all kinds of discrete dynamic systems, combining common abstraction principles from computational logic and discrete mathematics.

Each formal modeling approach focuses on a certain view towards systems, being declarative, functional, or operational. Some languages are particularly good in modeling data structures and the state of systems but are less supportive on the operational aspects. Some are low level, staying closer to code and the final implementation of systems and some are more formal and stay on the mathematical level. Among these formal methods, abstract state machines, while being primarily operational in nature, provide a good compromise between declarative, functional and operational views towards modeling distributed discrete dynamic systems. The emphasis on *freedom of abstraction* [25] in ASMs leads to intuitive yet accurate descriptions of systems which, thanks to the pseudo-code style of its language, are easily understandable by both domain experts and system designers. Since ASMs are in principle executable, the resulting models are validatable and possibly falsifiable by experiment. Finally, the well-defined notion of *step-wise refinement* in ASMs [18] bridges the gap between the abstract model and its final implementation.

1.2 Formal Language Semantics

Modeling languages are used to create a formulation of a system, based on one's understanding of that system or its requirements, so that it can be documented, communicated

[5] `http://www.petrinets.info`

[6] A *bipartite* graph is a graph that does not contain any odd-length cycle.

with peers and domain experts, and better yet, empirically validated if possible. Such a formulation needs to be clear, precise and comprehensive at a given level of abstraction. In order to achieve this, one needs to have a good understanding of the underlying modeling language used, which in turn requires a "good" description of the language.

A complete description of a modeling language covers three aspects of the language: *syntax*, *semantics*, and *pragmatics* [110]. The syntax is about the superficial form of the language constructs. It answers questions like, "is X a proper statement in this language?" The semantics is about the interpretation and the meaning of statements of the language. It answers questions like, "what does '`x := y + 1`' mean and what are its effects?" Finally, the pragmatics is about the use of such statements.

In this section, we focus on formal approaches towards semantic description of languages. Language descriptions used to be more informal (i.e., expressed in a narrative form using a natural language), since formal descriptions using rigorous notations are not easily understandable without special training. However, it is often difficult to precisely and clearly describe the semantics of languages using an informal language. Informal descriptions rely on a common understanding of the underlying informal language and are amenable to different interpretations which in case of modelling languages defies the purpose of having a clear, precise and comprehensive formulation. If we want completeness, consistency, precision, absence of ambiguity, and understandability, we have to look into formal descriptions.

Formal semantic specification of a language can serve many purposes [119, 110]:

1. *Reference for users:* A formal specification can serve as a reference for users of the language, providing a detailed and accurate description of the language, its meaning and its effects.

2. *Reference for implementations.* Those who implement tools for a language such as compilers, interpreters or debuggers, need to precisely know the details of the language and its semantics. Also, such specifications are needed if one wants to prove the correctness of language compilers or interpreters.

3. *Improved language design.* Formal specifications can expose irregularities and inconsistencies in language design and can guide language designers towards the design of better and cleaner languages.

4. *Standardization.* It is now generally accepted that formal specifications are necessary to have a successful language standardization process.

5. *Program/model verification.* To mathematically prove the correctness of models and programs, the properties of the underlying language constructs must be formally defined.

There are three main approaches to the semantic formalization of programming languages: operational, functional (denotational), and declarative (axiomatic). In the following sections, we briefly look into these approaches.

1.2.1 Operational Semantics

The essence of the operational approach is to explain the semantics of the language by defining an *abstract machine* with discrete "states" that takes the terms of the language as input [119, 110]. The machine interprets the programs of the language and performs the sequences of actions specified by the program by passing through a sequence of discrete states.

The abstract machine used in this approach itself usually comes with a formal definition, but the idea is that the code written for this machine is often clear, precise and simple, and much easier to understand than the language it is used to specify.

One of the oldest metalanguages used in such descriptions is the Vienna Definition Language, VDL for short, which was used for the definition of the PL/I language in 1969 [110, 119, 92]. A more recent example of the application of the operational method for formal semantics is the formal specification of the Java language and the Java Virtual Machine [118], using the abstract state machine framework, that aims to enable mathematical or computer-assisted verification as well as experimental validation of certain properties of Java. In the past two decades, abstract state machines have been used for semantic foundations of various industrial system design languages like the ITU-T standard for SDL [69, 45, 44, 83], the IEEE language VHDL [21, 20] and its successor SystemC [108], programming languages like C# [19] and Prolog [14, 15], and Web service description languages [55, 54].

1.2.2 Denotational Semantics

Denotational semantics, developed by Dana Scott and Christopher Strachey [115], is a formal method of describing the semantics of programming languages in form of abstract mathematical objects. As originally introduced by Scott and Strachey, denotational semantics provided the meaning, or *denotation*, of a program as a mathematical function that maps a program's input into output.

The denotational semantics approach is more abstract than the operational semantics, and it is mostly oriented toward language designers [115, 110]. The notion of "state" still exists in an abstract form but there is no explicit modeling of computational processes such as interpretation of programs. Abstract syntax of the subject language is modeled as a set of syntactic domains (such as variables, expressions and statements). The meaning of program statements are provided by semantic functions that map the statements into elements of semantic domains, elements of which are usually mathematical functions representing the *meanings* that can be assigned to the statements of the language.

The final piece in this puzzle is a semantic function that maps syntactic domains to semantic domains; i.e., assigns meanings to syntactic elements of a language. After defining the type (or the signature) of the semantic function, its detailed definition (i.e., the semantic specification of the language) is defined by a set of semantic equations, one for each summand of the syntactic domain.

Action Semantics

Peter D. Mosses introduced the *action semantics* method [106], developed based on the denotational semantics, to overcome some of the pragmatic problems he observed with application of denotational semantics to "realistic" programming languages; problems such as complexity of the semantic specifications, poor modifiability and extensibility [106, 107]. Hence, action semantics is developed with the goal of enriching denotational features with practically useful operational ones [16].

Like denotational semantics, action semantics involves defining semantic functions that map abstract syntax, in a compositional pattern, to semantic entities [106]. The main differences between these two approaches are the nature of *semantic entities* and the notation used to express them. As opposed to functions (or denotations) used in denotational semantics, there are three kinds of semantic entities used in action semantics: *actions*, *data*,

and *yielders*. Actions are dynamic computational entities, performance of which represents a computational behavior. Data items are static mathematical entities, representing pieces of information, and yielders represent unevaluated pieces of data.

1.2.3 Axiomatic Semantics

Axiomatic semantics, originally introduced by C. A. R. Hoare [78], is a method of specifying semantics of languages by formal statements about the *effect* of executing the language constructs. It is considered to be the most high-level approach to semantics specification.

The idea behind axiomatic semantics is that a semantic specification of a language is sufficiently defined if based on that specification one can prove any provable true statement (and no false statement) about the language [110]. Axiomatic semantics is more concerned with the general problem of program verification and synthesis (making it suitable for proving certain properties of programs written in that language) and is oriented toward language users. There is no explicit notion of a "state" or a machine; instead the focus is on defining an ideally minimal set of semantics constraints that must be satisfied by any "correct" implementation of the language. There is no indication of how such an implementation should be achieved.

There is no single standard metalanguage for defining axiomatic semantics of languages. The semantics is specified mostly in form of *assertions* (formulas in predicate logic) on the values of program variables or the relationship between those values. The class of assertions is extended with the *Hoare Triple* of the form:[7]

$$\{P\}\ S\ \{Q\}$$

where P and Q are precondition and postcondition assertions and S is a statement or construct in the subject language. The semantics of this assertion is: *if P holds before the execution of S, and if the execution of S successfully terminates, then Q holds after the execution of S* [78, 110]. It is important to note that if S does not terminate, then there is no "after" and Q can be any statement. Hence, the Hoare logic can only prove *partial correctness*, i.e. correctness subject to assumption of termination. Termination of S would have to be proved separately.

[7] This assertion is sometimes written as "$P\ \{S\}\ Q$", as it is originally introduced by Hoare [78].

An axiomatic specification of a programming language includes a number of *deduction rules* (or *rules of inference*) which allows deduction of the truth of assertions based on the truth value of other assertions. The general form of a deduction rule is

$$\frac{H_1, H_2, \ldots, H_n}{H}$$

in which $H1, \ldots, H_n$ and H are assertions and the interpretation is: *if* $H1, \ldots, H_n$ *are true, it may be deducted that* H *is true.* For example, Hoare's rule of composition

$$\frac{\{P\}S_1\{Q\}, \{Q\}S_2\{R\}}{\{P\}S_1; S_2\{R\}}$$

applies to sequential execution of two statements S_1 and S_2.

1.3 Towards a Comprehensive Framework

In light of such observations, a question naturally comes to mind: *what does it take to develop a comprehensive framework and tool environment for design and modeling of complex distributed systems and what features should such a framework provide?* Building on our experience with a broad scope of applications spanning web services architectures [56], computational criminology [51], maritime surveillance [58] and situation analysis [50], we believe that the following set of requirements should be satisfied by any such framework:

1. *Simple and concise specifications*

   Specifications written in such a framework should be simple and concise to be readable and understandable by both domain experts and system designers and to facilitate reasoning about the design and the communication of design concepts between those groups.

2. *Precise semantic foundation*

   The modeling language of such a framework should come with a precise semantic foundation as a prerequisite for analysis, validation and verification of the models.

3. *Freedom of abstraction*

   Such a framework should support writing of abstract and minimal specifications that express the original idea behind the designs of systems at the same levels of complexity and enable system designers to stress on the essential aspects of their design rather than encoding the insignificant details.

4. *Design exploration through fast prototyping*

   Exploring the problem space for the purpose of writing an *initial specification* requires a language that emphasizes freedom of experimentation by minimizing the need for encoding in mapping the problem space to a formal model. This can be achieved by

   - *reducing the cost of encoding domain concepts to language concepts* by providing a rich set of abstract data structures, various domain-specific concepts, and extensibility mechanisms for the tool environment and its language,
   - *avoiding early commitments* and encouraging rapid prototyping by supporting creation of abstract and untyped models that can later be refined into more concrete models.

5. *Refinement of models*

   Support for abstraction should be paired with a well-defined refinement technique that allows the system designer to cross levels of abstraction and link the models at different levels through incremental steps down to the final implementation (or the concrete model).

6. *Executability of specifications*

   Executability of even fairly abstract and incomplete models is important to allow experimental validation of the specifications at the early stages of design and to improve communication with the stake-holders during the requirements elicitation and analysis process.

7. *Support for distributed models (multi-agent systems)*

   It is only natural to expect a framework for design and modeling of distributed systems to explicitly support distributed and multi-agent design. This includes support for the definition of agent programs (or processes), inter-agent interaction mechanisms, and various scheduling policies.

8. *Non-determinism*

   Non-determinism is useful as a means of abstracting away from details of complicated and potentially deterministic algorithms. For example, non-deterministic descriptions can be used in high-level modeling of the behavior of the environment.

Considering these requirements, we argue that the ASM formalism properly matches our needs as the underlying formal framework for such a tool environment:

- Abstract state machine specifications are in fact rigorously-defined pseudo-code programs on abstract data structures [25]. As a result, they support writing of simple and concise specifications with a precise semantic foundation.
- ASM programs and the data structures can be fairly abstract[8] and yet ASM specifications are in principle executable.
- The ASM framework comes with a sound and powerful notion of step-wise refinement that helps the designer to structure the design of a system into appropriate abstraction levels and link those levels down to the concrete model (or code).
- The ASM formalism supports the design of distributed systems by providing two classes of synchronous and asynchronous multi-agent abstract state machines.
- ASM supports non-determinism in two forms: a *choose* construct that conveniently abstracts from the details of scheduling, and the notion of read-only *monitored functions* that are only updated by the environment of the system.

Looking at past experiences with ASM languages and modeling environments and considering the requirements listed above, we reason that a comprehensive ASM framework for design and analysis of distributed systems should:

1. come with a rich ASM language that supports both basic and distributed ASMs with non-determinism (see Chapter 2);
2. offer a formal (preferably operational) specification of its language and simulation engine that ensures
   - precise semantics,
   - preservation of pure ASM semantics, and
   - executability of the language;
3. ensure freedom of experimentation through extensibility of the language and its environment;
4. support interaction with the environment (e.g., external functions);

[8] In ASMs arbitrary structures can be used to reflect the underlying notion of state [25, P. 22].

5. be implemented as an *open framework* under an open source license[9] and using a platform-independent language and architecture so that it can be later modified or improved as needed by its users.

It would also be an advantage if such a framework provides a GUI (Graphical User Interface) for simulation and debugging. The graphical interface can organize the information relevant to state transitions into different views, visually highlight inconsistencies of the model, and give the user the ability to compare and contrast states and updates produced by different steps.

1.4 The CoreASM Modeling Environment

We take into account the requirements discussed above in the design and development of CoreASM to offer one instantiation of such a comprehensive framework for high-level design and analysis of distributed systems. In this section, we look into different aspects of design and implementation of CoreASM and address some of the challenges one may face during its development.

Formal Specification

There is no need to argue that the development of a reliable modeling framework for design and analysis of distributed systems has to start with a formal (read precise) specification of its language and tool architecture. Abstract state machines have been extensively used for semantic foundations of various programming and system design languages (see Chapter 2). While ASM specifications are primarily operational in nature, they provide a good compromise between declarative, functional and operational views toward modeling of languages and systems. Hence, it is only reasonable to use ASMs in formal modeling of the CoreASM language and its simulation environment (see Chapter 4).

We specify the CoreASM language (both its syntax and the corresponding semantics) through the specification of an interpreter (in form of an abstract state machine), therefore ensuring the executability of the language while providing its formal semantics. The design of the simulation engine and its architecture are specified using Control State ASMs [25], a

[9] http://www.opensource.org

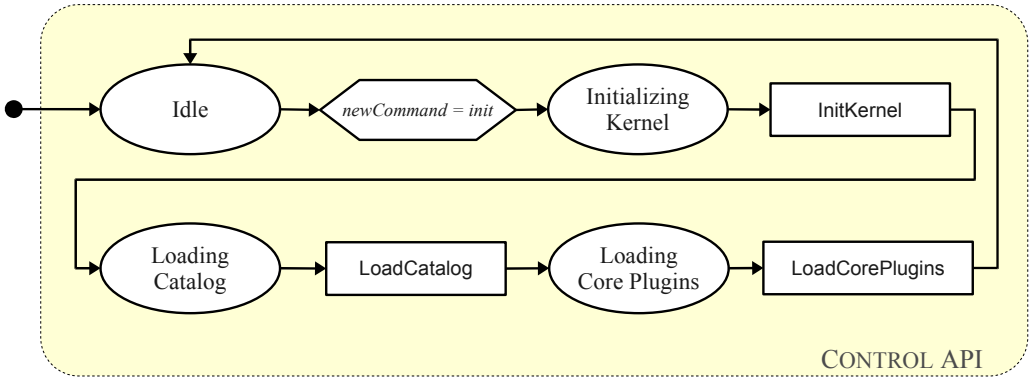

Figure 1.1: An Example of a Control State ASM

practical class of abstract state machines that have an easy-to-understand graphical representation (see Figure 1.1 for an example).

Extensible Architecture

In order to provide a rich ASM language that preserves pure ASM semantics and supports sequential and distributed ASMs with non-determinism, we closely follow the formal semantics and the definition of ASMs as provided by the ASM book [25]. However, this may not be enough. ASMs have been used in various domains, some of which required the introduction of special rule forms and data structures into ASMs. To follow the same spirit and to preserve this freedom of experimentation that comes with ASMs, the CoreASM language has to be easily extensible by third parties so that it can naturally fit into different application domains. In addition, to ensure freedom of experimentation, we would like to allow various modeling tools and environments to closely interact with the engine and also to let researchers experiment with variations to the engine's functionality. As a result, we propose a *plugin-based architecture* with a minimal kernel for the CoreASM language and modeling environment to offer the extensibility of both the language and its simulation engine. We start with a micro-kernel (the *core* of the language and its engine) that contains the bare essentials, that is, all that is needed to execute only the most basic ASM. We then implement most of the constructs of the language and the functionalities of the engine through plugins extending the kernel.

Language extensibility is not a new concept [117]. There are a number of programming

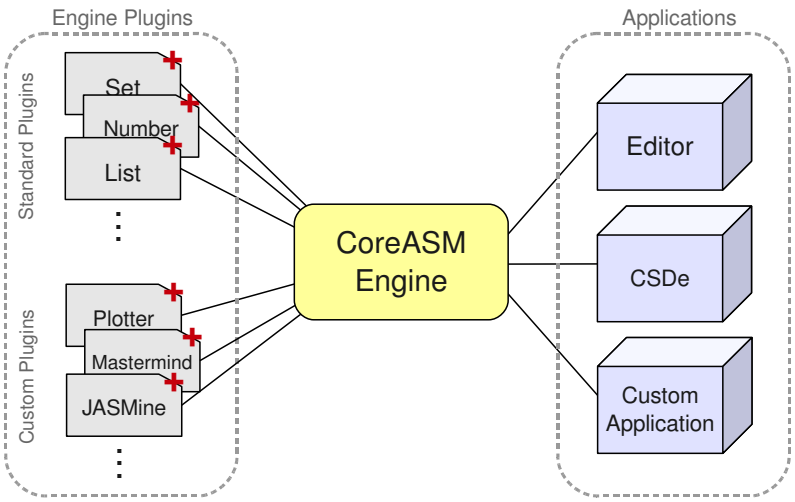

Figure 1.2: CoreASM Extensible Architecture

languages that support some form of extensibility from defining new macros to the definition of new syntactical structures. However, what we are suggesting here is the possibility of extending and modifying the syntax and semantics of the language, keeping only the bare essential parts of the ASM language as static. In order to achieve the this goal, CoreASM plugins should be able to extend the grammar of the core language by providing new grammar rules together with their semantics (see chapters 5 and 6). As a result, every time a CoreASM specification is being loaded, based on the set of plugins that the specification uses, the engine builds a language and a parser for that language to parse the specification. Since the set of all the possible plugins and their grammar rules is not known at the design time (which would otherwise defy the purpose of having a plugin-based architecture) one of the challenges would be to to equip the engine with a fast parser generator capable of generating parsers with look-ahead of more than one to allow the co-existence of more than one grammar rule starting with the same pattern.

Implementation

To facilitate the integration of CoreASM with other complementary tools such as symbolic model checking and automated test generation, the CoreASM engine should have a sophisticated and well defined interface to its environment which provides an API for various operations such as loading a CoreASM specification, starting an ASM run, or performing a

single execution step.

In order to have an open and platform-independent implementation of CoreASM, the whole framework is implemented in Java under an open source license (see Chapter 7). After considering various open source license models and looking at similar open source projects, we decided to make CoreASM source code available under the Academic Free License (AFL) version 3.0[10]. AFL 3.0 is an open source license with no reciprocal obligation to disclose source code; i.e., derivative works can be licensed under other licenses, and the source code of those derivative works need not be disclosed. Such a license provides a good compromise between the availability of the original source code in a free form and the existence of potentially proprietary editions and extensions in the industry.

1.5 Thesis Organization

The rest of this thesis is organized as follows. Chapter 2 provides an introduction to abstract state machines and uses an example to illustrate the application of ASMs in modeling industrial systems. An overview of the related work, mainly focusing on other ASM tools and modeling environments, closes our introductory material in Chapter 3.

Chapter 4 opens the second part of the thesis with an overview of the architecture of CoreASM and its main components. An in-depth description of the kernel of CoreASM is then provided in Chapter 5, followed by the specification of currently available CoreASM plugins, presented in Chapter 6, that extend the functionalities of the CoreASM kernel and gradually form the comprehensive CoreASM framework.

The concluding part of this thesis begins with Chapter 7 explaining how the specification and design of the CoreASM kernel and its plugins are implemented in form of a Java application and plenty of complementary modules. Chapter 8 examines the application of CoreASM in high-level design and analysis of distributed systems, and Chapter 9 concludes the thesis by addressing the significance of the work and laying out the subjects of future improvements.

[10] http://www.opensource.org/licenses/afl-3.0.php

Chapter 2

Abstract State Machines

Abstract State Machines (ASMs), originally known as *Evolving Algebras*, were first introduced by Yuri Gurevich [73, 74] as a versatile mathematical method of modeling discrete dynamic systems with the goal of bridging the gap between computation models and specification methods. ASMs combine two well-known and fundamental concepts of *transition systems*, to model the dynamic aspects of a system, and *abstract states*, to model the static aspects at any desired level of abstraction. Egon Börger [25] further developed ASMs into a *systems engineering* method that guides the development of software and embedded hardware-software systems from requirements capture to their implementation.

Today, ASMs are well known for their versatility in computational and mathematical modeling of architectures, languages, protocols and virtually all kinds of sequential, parallel and distributed systems with an orientation towards practical applications. The particular strength of this approach is the flexibility and universality it provides as a mathematical framework for semantic modeling of functional requirements in terms of abstract machine models and their runs. Widely recognized applications of ASMs include semantic foundations of industrial system design languages like the ITU-T standard for SDL [69, 45, 44, 83], the IEEE language VHDL [21, 20] and its successor SystemC [108], programming languages like JAVA [118, 24], C# [19] and Prolog [14, 15], Web service description languages [55, 54, 53], communication architectures[70, 71], embedded control systems [23, 8, 22], et cetera.[1]

In this chapter we briefly recall the basic notions of ASMs as defined in [25] and we

[1] See also the ASM website at `www.asmcenter.org` and the overview in [25].

use an example to illustrate the application of ASMs with CoreASM in modeling industrial systems.

2.1 Basic ASMs

The original notion of ASMs, or *basic ASMs*, was defined to formalize simultaneous parallel actions of a single computing agent. This notion was later generalized to capture the formalization of multiple agents acting and interacting in an asynchronous manner [25]. In this section, we focus on basic ASMs. *Multi-agent ASMs* or *Distributed ASMs* are explored in the next section.

2.1.1 Basic Definition

A basic ASM M is a tuple of the form $(\Sigma,\ \mathcal{I},\ \mathcal{R},\ P_M)$ where:

- Σ is a signature; i.e., a finite set of function names f where each function has an *arity*, which is the number of arguments that function takes. Nullary functions, those with arity of zero, are called *constants*. The constants *true*, *false*, and *undef* (representing the "undefined" value) are always defined.

- $\mathcal{I}$ is a set of initial states for signature Σ. A state $\mathfrak{A}$ for Σ is a non-empty set X (the *superuniverse* of $\mathfrak{A}$) together with an interpretation $f^{\mathfrak{A}}$ for each function name f in Σ such that:

  - if f is an n-ary function name, then $f^{\mathfrak{A}} : X^n \mapsto X$, and
  - if c is a constant in Σ, then $c^{\mathfrak{A}} \in X$.

  Functions can be *static* or *dynamic*. Values of dynamic functions can change from state to state.

- $\mathcal{R}$ is a set of rule declarations. In a given state, evaluation of a rule $r \in \mathcal{R}$ produces an *update set* of updates of the form (l, v) where:

  - l is a *location*. A location l in state $\mathfrak{A}$ is a pair $(f, \langle a_1, \ldots, a_n\rangle)$ where f is an n-ary function name in Σ and $a_1, \ldots, a_n$ are values from superuniverse X (i.e., $\forall_{i\in\{1,\ldots,n\}} a_i \in X$). The contents of a location l in $\mathfrak{A}$ is $f^{\mathfrak{A}}(a_1, \ldots, a_n)$.
  - v is a value of superuniverse X.

The meaning of an update (l, v) is that the content of location l has to be changed to the value v.

- $P_M \in \mathcal{R}$ is a distinguished rule of arity zero (no free variables), called the *main rule* or the *Program* of machine M.

The superuniverse X is usually divided into smaller *universes* modeled by their characteristic functions (unary relations). If D is a universe, then the set of all elements of D is defined as $\{d \mid D(d) = \mathit{true}\}$.

2.1.2 State Transitions

ASM specifications describe how the state of the specified system evolves in time. A computation of M, starting with a given initial state $S_0 \in \mathcal{I}$, results in a finite or infinite sequence of consecutive state transitions of the form

$$S_0 \xrightarrow{\Delta_{S_0}} S_1 \xrightarrow{\Delta_{S_1}} S_2 \xrightarrow{\Delta_{S_2}} \cdots,$$

such that S_{i+1} is obtained from S_i, for $i \geq 0$, by *firing* Δ_{S_i} on S_i, where Δ_{S_i} denotes a consistent finite set of updates computed by evaluating P_M over S_i.

An update set is called *consistent* if it does not have clashing updates that attempt to assign different values to the same location. The result of firing a consistent update set Δ_{S_i} on S_i is a new state S_{i+1} with the same superuniverse as S_i, such that for every location l of S_i we have:

$$S_{i+1}(l) = \begin{cases} v, & \text{if } (l, v) \in \Delta_{S_i} \\ S_i(l), & \text{otherwise.} \end{cases}$$

2.1.3 Transition Rules

The program P_M of an ASM M is defined by an ASM transition rule.[2] Basic transition rules are as follows:

1. *Skip rule:* **skip**

   Does nothing and evaluates into an empty update set.

[2] This is a pragmatically generalized definition based on the original definition of an ASM program by [25] which defines an ASM [program] as a set of guarded transition rules.

2. *Update rule:* $f(a_1, \ldots, a_n) := t$
   Updates the value of $f(a_1, \ldots, a_n)$ to t. It evaluates into an update set of the form $\{(f(a_1, \ldots, a_n), t^{\mathfrak{A}})\}$ where $\mathfrak{A}$ is the current state of the machine and $t^{\mathfrak{A}}$ is the value of t in $\mathfrak{A}$.

3. *Block rule:* P **par** Q
   Evaluates rules P and Q in parallel and the result is the union of the update sets computed by P and Q.

4. *Conditional rule:* **if** ϕ **then** P **else** Q
   If ϕ is true, this rule executes P, otherwise executes Q.

5. *Let rule:* **let** $x = t$ **in** P
   Assigns the value of t to x and executes P. The resulting update set is the update set produced by P.

6. *Forall rule:* **forall** x **with** ϕ **do** P
   Executes P in parallel for every x that satisfies ϕ. The resulting update set is the union of all the update set produced by parallel execution of P over different values of x.

7. *Choose rule:* **choose** x **with** ϕ **do** P **ifnone** Q
   Non-deterministically (unless otherwise specified) chooses x satisfying ϕ and executes P. If no such x exists, it executes Q.

8. *Sequence rule:* P **seq** Q
   Execute P, if the update set produced by P is consistent, then execute Q in a state which the updates of P are applied. The resulting update set U (based on U_P and U_Q update sets of P and Q) is

$$U = \begin{cases} \{(l, v) \in U_P \mid l \notin \mathit{locations}(U_Q)\} \cup U_Q, & \text{if } U_P \text{ is consistent;} \\ U_P, & \text{otherwise.} \end{cases}$$

9. *Call rule:* $\mathsf{R}(a_1, \ldots, a_n)$
   Execute the previously defined transition rule R with the given parameters. Parameters are passed in a *call-by-name* fashion; i.e., they are passed unevaluated. ASM

transition rules can be defined using the expression

$$\mathsf{R}(x_1, \ldots, x_n) = P$$

where R is the name of the new rule, P is a transition rule and the free variables of P are included in $x_1, \ldots, x_n$.

2.1.4 Interaction with Environment

M interacts with a given operational environment—the part of the external world visible to M—through actions and events as observable at external interfaces, formally represented by externally controlled functions. Intuitively, such functions are manipulated by the external world rather than M itself. Of particular interest are *monitored functions*. Such functions change their values dynamically over runs of M, although they cannot be updated internally by agents of M. A typical example is the abstract representation of global system time. In a given state S of M, the global time (e.g., as measured by some external clock) is given by a monitored nullary function *now*, taking values in a linearly ordered domain TIME $\subseteq$ REAL. Values of *now* increase monotonicly over runs of M.

2.2 Multi-Agent ASMs

Basic ASMs are extended to capture the formalization of multiple agents acting and interacting in an asynchronous manner [25].[3]

An asynchronous multi-agent ASM (or DASM for Distributed ASM) M^D is defined by a dynamic set AGENT of computational *agents* each executing its ASM. This set may change dynamically over runs of M^D, as required to model a varying number of computational resources. Agents of M^D normally interact with one another, and typically also with the operational environment of M^D, by reading and writing shared locations of a global machine state.[4]

A DASM M^D performs a computation step whenever one of its agents performs a computation step. In general, one or more agents may participate in the same computation step

[3] A synchronous version of multi-agent ASMs also exists [25, Sec. 5], in which a set of agents execute their own programs in parallel, synchronized by an implicit global system clock. Since asynchronous ASMs are more general, we will not further explore synchronous ASMs in this survey.

[4] In principle, one may also compose a DASM of a number of agents, each operating on a part of the state that is disjoint from the view of all the other agents, so that each agent has its own private state.

of M^D. A single computation step of an individual agent is called a *move*. In this model, moves are atomic. Naturally, conflicting moves must be ordered so that they do not occur in the same step of M^D.

A partially ordered run ρ of M^D is given by a triple (Λ, A, σ) satisfying the following four conditions (adopted from [74, Sec. 6.5]):[5]

1. Λ is a partially ordered set of moves, where each move has only finitely many predecessors.

2. A is a function on Λ associating agents to moves such that the moves of any single agent of M are linearly ordered.

3. σ assigns a state of M to each initial segment X of Λ, where $\sigma(X)$ is the result of performing all moves in X.

4. *Coherence condition*: If x is a maximal element in a finite initial segment X of Λ and $Y = X - \{x\}$, then $A(x)$ is an agent in $\sigma(Y)$ and $\sigma(X)$ is obtained from $\sigma(Y)$ by firing $A(x)$ at $\sigma(Y)$.

A partially ordered run defines a class of admissible runs of M^D rather than a particular run. In general, it may require more than one (even infinitely many) partially ordered run to capture all admissible runs of M^D. From the coherence condition it follows that all *linearizations* of the same finite initial segment of a run of M^D have the same final state.[6] The implication of the partially-ordered-run semantics is illustrated by means of a simple but meaningful example.

Example: Door and Window Manager Assume two propositional variables, *door* and *window*, where $door = true$ means that 'the door is open' and $window = true$ means that 'the window is open'. There are two distinct agents: a door-manager d and a window-manager w.

[5] Here we recall our notes from [49].

[6] Intuitively, a finite initial segment of a partially ordered run ρ is a finite subset of Λ corresponding to a (finite) prefix of ρ.

Door/Window Managers

DoorManager ≡
  if ¬*window* **then** *door* := *true* // move x

WindowManager ≡
  if ¬*door* **then** *window* := *true* // move y

Initially (in state S_0) both the door and the window are closed. Then there are only two possible runs, and in each run only one of the agents makes a move.

We cannot have $x < y$ because w is disabled in the state S_x obtained from S_0 by performing x. Also, we cannot have $y < x$ because d is disabled in the state S_y obtained from S_0 by performing y. Finally, we cannot have a run where x and y are incomparable, that is neither $x < y$ nor $y < x$. By the coherence condition, the final state $S_{x,y}$ of such a run would be obtained from either S_x by performing y or from S_y by performing x; either case is impossible.

2.3 Control State ASMs

In this section we briefly look into *control state ASMs*, a frequently used class of ASMs that represents a normal form of synchronous UML activity diagrams. This particular class of ASMs is expressive enough to model many classical automata such as various extensions of finite state machines, timed automata, push-down automata, etc. It extends finite state machines by synchronous parallelism and by the possibility to also manipulate data [25].

A control state ASM is an ASM whose rules are all of the form presented in Figure 2.1.[7] Such a control state ASM can be formulated in textual form by a parallel composition of Finite State Machine (FSM) rules, where each FSM rule is defined as:

FSM(*i*, **if** *cond* **then** *rule*, *j*) ≡
  if *ctl_state* = *i* **and** *cond* **then**
    rule
    ctl_state := *j*

Thus, the control state ASM of Figure 2.1 can be formulated as a parallel composition of the following FSM rules:

[7] See [25, Sec. 2.2.6]

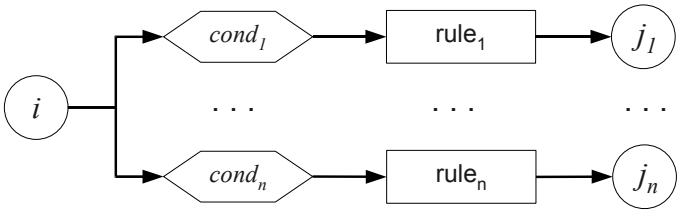

Figure 2.1: Control State ASMs

FSM$(i, \textbf{if } cond_1 \textbf{ then } rule_1, j_1)$
FSM$(i, \textbf{if } cond_2 \textbf{ then } rule_2, j_2)$
...
FSM$(i, \textbf{if } cond_n \textbf{ then } rule_n, j_n)$

Since control state ASMs can be presented in graphical form with a precise semantics, they are a good candidate for documenting functional requirements and modeling of functional aspects of systems at the early stages of design and development when proper communication of the requirements and the abstract model plays a key role.

2.4 Similar Approaches in Computational Logic

There are similar approaches to formal modeling of systems, the B method (see Section 1.1) being one of the most popular ones. The idea of modeling states of the system as algebraic structures have been practiced earlier in various forms of Algebraic Specifications (see Section 1.1). However, in the area of computational logic where ASMs are coming from, one of the closest formal modeling approaches to ASMs with respect to their view towards systems, runs, and distributed computation seems to be the *interpreted systems* approach.

Fagin et al. introduced the notion of *interpreted systems* [46] as a formal semantic framework for reasoning about knowledge and uncertainty in multiagent systems. It is interesting to observe the similarities between abstract state machines and the multi-agent modeling framework of interpreted systems, specially in capturing the notions of agents, concurrency, runs, update actions, and programs. Here, we briefly recall [50] the basic notions and definitions of the underlying systems modeling framework of interpreted systems.

2.4.1 Runs and Systems

According to [46], a multiagent system can be conceptually divided into two components: the agents $A = \{a_1, \ldots, a_n\}$ and the environment e, which can be viewed as a special agent. The *global state* of the system with n agents is defined to be an $(n+1)$-tuple $(s_e, s_1, \ldots, s_n)$, where s_e is the state of the environment and s_i is the local state of agent i. The set of all global states of the system is defined as $\mathcal{G} : L_e \times L_1 \times \ldots \times L_n$, where L_e is the set of possible states for the environment and L_i is the set of all possible local states of agent i. To model the changes of the system's global state in time, the notion of *run* is introduced as a function from time to global states $\mathcal{G}$, with the assumption that time ranges over natural numbers. A system can have many possible runs. The initial global state of a system with a possible run r is $r(0)$. A pair (r, m) consisting of a run r and a time m is referred to as a *point* in run r.

A *system* $\mathcal{R}$ *over* $\mathcal{G}$ is defined as a set of runs over $\mathcal{G}$.

2.4.2 Actions, Protocols, and Programs

A *round* takes place between two points in a run, and a round m in run r is defined to take place between points $r(m-1)$ and $r(m)$. Agents and the environment change the global state by performing *actions* in rounds. Let ACT_i be the set of actions that can be performed by agent i, and let ACT_e be the set of actions that can be performed by the environment. A *joint action* is a tuple $(a_e, a_1, \ldots, a_n)$ of actions performed by the environment and the set of agents, where $a_e \in ACT_e$ and $a_i \in ACT_i$ for i in $1 \ldots n$.

Joint actions cause the system to change its global state and the change is modeled by a *global state transformer* function $\mathcal{T} : \mathcal{G} \mapsto \mathcal{G}$ that is associated to each joint action $(a_e, a_1, \ldots, a_n)$. A *transition function* τ is a mapping that associates a global transformer with each joint action. It is required that $\tau(a_e, a_1, \ldots, a_n)(s_e, s_1, \ldots, s_n)$ be defined for each joint action $(a_e, a_1, \ldots, a_n)$ and each global state $(s_e, s_1, \ldots, s_n)$.

Agents perform actions according to some *protocol*, which is a rule for selecting actions. A protocol P_i for the agent i is formally defined as $P_i : L_i \mapsto \mathcal{P}(ACT_i)\backslash\{\emptyset\}$. A protocol P_i is *deterministic* if $\forall s_i \in L_i \;\; |P_i(s_i)| = 1$. In a similar fashion, a protocol P_e for the environment is defined as a function from L_e to nonempty subsets of ACT_e. A *joint protocol* P is a tuple $(P_1, \ldots, P_n)$ consisting of all the protocols P_i, for each of the agents $i = 1, \ldots, n$. Note that the environment's protocol P_e is not included in the joint protocol. The protocol of the

environment is usually supposed to be given and P and P_e can be viewed as the strategies of opposing players.

A context γ is defined as a tuple $(P_e, \mathcal{G}_0, \tau, \Psi)$, where P_e is a protocol for the environment, $\mathcal{G}_0$ is a nonempty subset of $\mathcal{G}$ describing the initial system state, τ is a transition function and Ψ is an *admissibility* condition on runs specifying which runs are "acceptable". Formally, Ψ is a set of runs; $r \in \Psi$ if r satisfies the condition Ψ. In practice, Ψ can be used to shrink down the system or to model fairness conditions. The combination of a context γ and a joint protocol P for the agents uniquely determines a set of runs.

Protocols are typically described by means of *programs* written in some programming language. A *standard program* for agent i is a statement of the form

```
case of
  if t_1 do a_1
  if t_2 do a_2
  ...
end case
```

where the t_j's are standard tests for agent i and the a_j's are actions of agent i (i.e., $a_j \in ACT_i$).

2.5 The Railroad Crossing Example

This section borrows the Railroad Crossing example of [25, Sec. 5.2.2] and offers a CoreASM model of the example to illustrate the application of CoreASM (and ASM in general) in modeling industrial systems.

A system controls a gate at a railroad crossing. There are multiple tracks on which trains can travel in both directions. There are sensors on the tracks that can detect if a train is *coming* or if it is currently *crossing*. The gate is controlled by two signals *open* and *close*. The purpose of the system is to keep the gate closed if a train is crossing (safety) and to keep it open otherwise (liveness).

2.5.1 The Abstract Model

We start our model by defining the universe of TRACK, initially set to include two tracks $track_1$ and $track_2$. We model the semantics of sensor values by defining a universe of

TRACKSTATUS; since the set of values are limited and known at the beginning, we model this universe as an enumerated universe. We also define an enumerated universe GATESTATE to capture two possible states of the gate: *opened* and *closed*.

```
universe Track = {track1, track2}
enum TrackStatus = {empty, coming, crossing}
enum GateState = {opened, closed}
```

The following function, *trackStatus*, holds the status of each track. Since there is only one gate in our system, a nullary function *gateState* is defined to keep the current state of the gate:

```
function trackStatus : Track -> TrackStatus
function gateState : -> GateState
```

The sensors are arranged such that when a train is detected as *coming*, it takes at least d_{min} seconds for it to arrive at the crossing. The gate takes d_{close} seconds to be closed and d_{open} to get opened. Thus, to keep the gate open as much as possible, if we detect a train coming we have $WaitTime = d_{min} - d_{close}$ seconds to start closing the gate. Hence, there is an implicit *deadline* associated to every track t, indicating the maximum time we have (with regard to track t) in order to safely close the gate.

```
function deadline : Track -> TIME
derived waitTime = dmin - dclose
```

The following nullary function *gateSignal*, controlled by the track control program, signals the opening or closing of the gate.

```
enum GateSignal = {open, close}
function gateSignal : -> GateSignal
```

The Rail Road Crossing ASM consists of two basic ASMs, TrackControl and GateControl, respectively controlling the tracks (sending signals to the gate controller) and maintaining the state of the gate (opening or closing the gate in response to gate signals). We assume that the environment sets the value of the function *trackStatus* based on the track sensors data.

The track control program TrackControl is a parallel combination of two main rules: 1) closing the gate if needed; i.e., for all tracks, calculating new deadlines, sending a close

signal if needed, and clearing passed deadlines; 2) opening the gate if it is safe to do so. The program is defined as follows:[8]

```
rule TrackControl = {
    forall t in Track do {
        SetDeadline(t)
        SignalClose(t)
        ClearDeadline(t)
    }
    SignalOpen
}
```

where we have

```
rule SetDeadline(x) =
    if trackStatus(x) = coming and deadline(x) = infinity then
        deadline(x) := now + waitTime

rule SignalClose(x) =
    if now >= deadline(x) and now <= deadline(x) + 1000 then
        gateSignal := close

rule ClearDeadline(x) =
    if trackStatus(x) = empty and deadline(x) < infinity then
        deadline(x) := infinity

rule SignalOpen =
    if gateSignal = close and safeToOpen then
        gateSignal := open
```

The predicate *safeToOpen*, used in the SignalOpen rule, can be defined as follows

$$safeToOpen \equiv \forall t \in \textsc{Track}\ trackStatus = empty \vee deadline(t) > now + d_{open}$$

which is defined in CoreASM as

```
derived safeToOpen = forall t in Track holds
                     trackStatus(t) = empty or deadline(t) > (now + dopen)
```

The gate control program simply responds to gate signals by changing the state of the gate:

[8] In CoreASM, curly braces {} can be used to define parallel rule blocks.

```
rule GateControl = {
    if gateSignal = open and gateState = closed then gateState := opened
    if gateSignal = close and gateState = opened then gateState := closed
}
```

2.5.2 The Executable Model

In order to have a meaningful execution of the model, we need to define the initial state of the system and simulate the behavior of the environment. So far we have defined two parallel ASM agents to model track and gate controllers. In this section we add two more agents to our model: an *Environment* agent to model the behavior of the environment and an *Observer* agent to observe the statuses of tracks and the gate and to provide a nicely formatted output throughout the simulation.[9] So, the universe of agents will be defined as:

```
universe Agents = {trackController, gateController, observer, environment}
```

The Environment

The environment agent simulates trains crossing over the tracks in a non-deterministic fashion. If a train is detected as *coming* on a track, we have d_{min} time before it crosses the intersection. Every train takes a certain time to pass the crossing; when that time is reached, the environment sets the track status back to *empty*. The following rule offers one possible definition of such an environment:

```
rule EnvironmentProgram =
    choose t in Track do {
        if trackStatus(t) = empty then
            if random < 0.05 then {
                trackStatus(t) := coming
                passingTime(t) := now + dmin
            }
        if trackStatus(t) = coming then
            if passingTime(t) < now then {
                trackStatus(t) := crossing
                passingTime(t) := now + 4000
            }
```

[9] However, we do not necessarily need to define these two agents in CoreASM. The environment can be modeled by monitored functions reading input from the user, and the printout can be generated using the Observer plugin presented in Section 6.4.4.

```
            if trackStatus(t) = crossing then
                if passingTime(t) < now then
                    trackStatus(t) := empty
    }
```

The Observer

The observer agent simply prints out the current state of the system. The following observer program prints out the current time, the statuses of all tracks, and finally the state of the gate. To keep the output lines in order, we enclose the print rules in a sequence block.

```
rule ObserverProgram =
    seqblock
        print "Time:  " + (( now - startTime) / 1000) + " seconds"
        forall t in Track do
            print "Track " + t + " is " + trackStatus(t)
        print "Gate is " + gateState
        print ""
    endseqblock
```

The Initial State

In CoreASM, the initial state of the system can be defined in an operational form using an *init* rule. The engine starts the execution of specifications by creating an init agent and assigning the init rule as the program of that agent (see Section 4.2). When the initial state is set up, the init agent can be de-activated by setting its program to *undef* or removing it from the universe of agents.

In our example, we assume that initially the gate is open, all the tracks are *empty* and track deadlines are set to positive infinity. The init rule, defined below, sets the initial values of functions and assigns the programs of the agents.

```
init InitRule

rule InitRule = {
    forall t in Track do {
        trackStatus(t) := empty
        deadline(t) := infinity
    }
    gateState:= opened
    dmin:= 5000
```

```
    dmax:= 10000
    dopen:= 2000
    dclose:= 2000
    startTime:= now

    program(trackController) := @TrackControl
    program(gateController) := @GateControl
    program(observer) := @ObserverProgram
    program(environment) := @EnvironmentProgram
    program(self) := undef
}
```

The Simulation

Finally, we have everything in place to execute the model in CoreASM and validate the behavior of the gate controller (see Appendix B.1 for the full specification). The execution provides a printout of the states of the system. The output shows that the controller keeps the gate open while there is no train on the tracks and keeps it closed as long as there is at least one train crossing the intersection. Figure 2.2 shows parts of the output of one particular run of the system. As a result of the non-deterministic behavior of the environment, different runs of the model most likely provide different outputs.

It is worth to emphasize that although the ability to execute the model and to observe its behavior enables us to validate the model by experiment, satisfying results of such experiments by no means guarantee the "correctness" of the model. Section 7.3.2 offers a brief discussion on this subject.

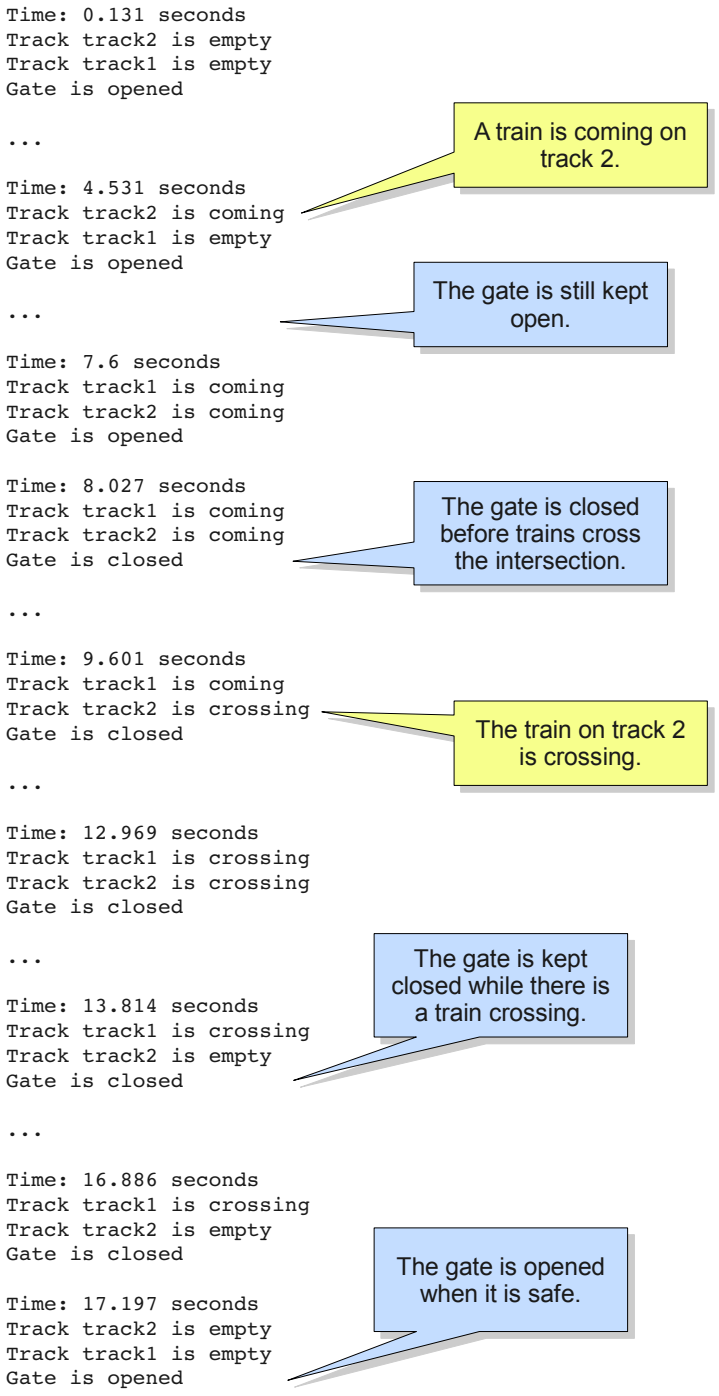

Figure 2.2: Output of the Railroad Crossing Example in CoreASM

Chapter 3

Related Work

Machine assistance plays an increasingly important role in making practical systems design feasible. Specifically, model-based systems engineering demands for abstract executable specifications as a basis for design exploration and experimental validation through simulation and testing. Thus, it is not surprising that there is a considerable variety of executable ASM languages that have been developed over the years.

The first generation of tools for running ASM models on real machines goes back to Jim Huggins' interpreter written in C [75, 81] and, even further back, to the Prolog-based interpreter by Angelica Kappel [87]. Other interpreters and compilers followed: the lean *EA* compiler [7] from Karlsruhe University, the *scheme*-interpreter [41] from Oslo University, and an experimental EA-to-C++ compiler developed at Paderborn University. Besides practical work on ASM tools, conceptual frameworks for more systematic implementations were developed. The work on the *evolving algebra abstract machine (EAM)* [39], an abstract formal definition of a universal ASM for executing ASM models, contributed to a considerably improved understanding of fundamental aspects of making ASMs executable.

Based on such experience, a second generation of more mature ASM tools and tool environments was developed: *AsmL* (ASM Language) [101] and the *Xasm (Extensible ASM) language* [4, 5] are both based on compilers, while the *ASM Workbench* [38], *AsmGofer* [113], and *Asmeta* [60] provide ASM interpreters.

All the above languages build on predefined type concepts rather than the untyped language underlying the theoretical model of ASMs. The most prominent of these languages are Asmeta and AsmL. The Asmeta language, called AsmetaL, implements all the constructs of

basic, structured, and multi-agent ASMs as defined in [25], but it is a fully typed ASM language with limited extensibility features. AsmL is a strongly typed language based on the concepts of ASMs but also incorporates numerous object-oriented features and constructs for rapid prototyping of component-oriented software, thus departing in that respect from the theoretical model of ASMs; rather it comes with the richness of a fully fledged programming language. Most of these languages do not provide a run-time system supporting the execution of distributed ASM models[1]; only Xasm (and Asmeta in a limited form) is designed for systematic language extensions; however, the Xasm language itself diverts from the original definition of ASMs and seems closer to a programming language.

The rest of this chapter reviews some of the more common and well-known ASM tools and languages (sections 3.1 to 3.6) and compares their features and shortcomings (see Table 3.5). Section 3.7 concludes the chapter with an overview of alternative tools in other state-based modeling languages.

3.1 The Dynamic Algebra Specification Language

In 1993, inspired by the work of Egon Börger on a dynamic algebra specification of full Prolog [14, 15], Angelica M. Kappel published a paper on the general concept of implementing dynamic algebras [87]. She defined a concrete language for dynamic algebra specifications, called DASL, and presented the design of an abstract algebraic target machine, called ALMA, specially tailored for dynamic algebra computations.

ALMA is a single-sorted abstract machine that provides three kinds of control statements: a simple statement and two conditional statements (*if-then* and *case* statements). An ALMA program is given by a decision tree in which the leaf nodes are either *update* or *error* nodes. An ALMA computation is the execution of a simple statement followed by a walk through the decision tree. The ALMA abstract machine and a compiler that translates DASL specifications to ALMA programs are both implemented in Prolog. The user interface is basically the Prolog environment.

In DASL, the user can explicitly create the initial state using the *start* rule. Regular transition rules are defined in terms of conditional statements and they are evaluated in every step of the simulation. To distinguish between error states and regular termination of

[1] Only Asmeta and AsmGofer provide some sort of support for the execution of distributed ASMs.

the machine, the set of regular final states must be explicitly defined in the specification.

DASL is a nice and clean specification language but it implements only a small subset of basic deterministic ASMs.

3.2 ASM Gofer

The *AsmGofer* system, designed and developed by Joachim Schmid [113], provides an ASM interpreter embedded in the functional programming language Gofer, a subset of the Haskell programming language. This interpreter has been used in a number of applications such as *Java and the Java Virtual Machine* [118], the *Light Control Case Study* [23], and *Simulating UML Statecharts* [37].

AsmGofer is in fact a conservative extension of Gofer adding the notions of state and parallel updates into Gofer. An AsmGofer program, or "script", is a collection of signatures, rules, functions and data structures that can appear in any order. Although AsmGofer is strongly typed, signatures are not mandatory.

Since ASM is a state-based modeling framework, every ASM update has an effect on the global state. This state-based view makes it challenging to implement an ASM language embedded in a pure functional programming language like Gofer that does not support side effects. To support the notions of 'state' and 'update', AsmGofer modifies the evaluation machine in Gofer run-time system and utilizes the *IO actions* in Gofer that are used for input-output operations [112]. In order to not change the Gofer syntax, ASM features are represented as expressions.

AsmGofer supports both classes of parallel and distributed ASMs. Support for distributed ASMs is provided by a special function, called *multi*, that gets a set of agents and non-deterministically chooses a subset of those agents in every step and executes their corresponding rules in parallel. It is important to note that this function never chooses a subset of agents that produce an inconsistent update set, since that would result in an invalid ASM run.

AsmGofer also support automatic GUI generation which is quite helpful in debugging and validation of specifications. Since monitored functions are not implemented, the GUI cannot be used for getting input from the environment or setting up initial values of functions.

3.3 XASM

One of the more advanced ASM simulators available is the XASM (*eXtensible ASM*) compiler and toolset developed by Matthias Anlauff [5]. The goal of the XASM project is to provide support for using ASMs as a programming language for producing efficient and reusable programs. As such, XASM focuses on the generation of efficient executable programs simulating the run of abstract state machines [4].

The XASM language supports all the transition rules defined in [74]. In order to simulate an ASM specification, XASM source files are translated into C source code by the XASM-compiler, which is then linked to the runtime system and optionally the user defined C functions. XASM introduces the concept of *component* (as defined in [121]) to ASMs by adding modularization constructs to its ASM language. Components in XASM programs can be reused either as a sub-machines contributing to the computation steps of the parent machine, or as computational functions modeled as independent machines with internal computation steps. Every XASM component, defined as an ASM, can provide a list of functions that it requires to *access* or can potentially *update*.

The XASM language supports interaction with external C programs in two ways. Specifically defined external C functions can be used in XASM specifications; however, the arguments and return values of C functions can only be of a specific C-type that represents elements of the super-universe in XASM. Alternatively, XASM-programs can be embedded in C-applications. An XASM specification, if compiled properly, can be included in C program and called as C function. Newer versions of XASM [5] support interaction with Java classes but the support is only limited to invoking Java object constructors.

XASM also provides support for pattern matching by introducing a pattern matching operator on strings that matches the left operand (as string data) with the right operand (as a pattern). The language also supports grammar definitions, which is very useful when ASM specifications are used to define programming language semantics.

The XASM compiler comes with a runtime system and a graphical debugger and animation interface that facilitates debugging and experimental validation of ASM specifications.

The XASM project is not maintained anymore.

3.4 The ASM Workbench

The ASM Workbench [35, 38], designed and developed by Giuseppe Del Castillo, is one of the most comprehensive modeling environments developed for ASMs. Realizing the shortcomings of available ASM tools of the time, such as incompatibility with other ASM tools and lack of support for formal analysis, Del Castillo started the ASM Workbench project to develop an open and extensible tool environment for ASMs. It was intended as a basis for the development of further ASM tools and complementary to other modeling tools [35].

The ASM Workbench is designed to be extensible. The main characteristics of its architecture are its *kernel* and a set of *exchange formats* that allow the kernel to be extended. The kernel of ASM Workbench is a set of program modules implemented in the functional language Standard ML [104]. It consists of a collection of data structures representing syntactic and semantic objects of ASM specifications, and a collection of functions to process ASM objects [36]. New functionalities can be implemented and added as extensions of the ASM Workbench either by *tight coupling* (writing other ML programs that use the functionalities of the kernel) or by *loose coupling* (writing other programs in any language that communicate with the ASM Workbench kernel, as a separate process, according to the exchange format conventions).

The ASM Workbench environment includes a type-checker, an interpreter, a graphical user interface for visualization and debugging, and an interface to the symbolic model checker SMV [88] (for a particular class of ASMs) based on transformation of ASMs into the SMV language.

The ASM language as defined in [74] focuses on specification of transition rules but but does not include any language constructs for specifying data structures, functions, and constraints. ASM Workbench introduces an ASM-based Specification Language (ASM-SL) that extends the basic language of ASM [74] by providing means to define the structure of the state, a simple and flexible type system, mechanisms to define interfaces to the external environment, and support for modularization.

The ASM-SL language is typed, adopting the type system of Standard ML. The argument for a typed language is clarity of specifications and improved error detection. It includes a set of predefined types, generic mathematical structures, and a few powerful constructions to define new types.

3.5 AsmL and Spec Explorer

On of the most prominent ASM languages is AsmL [101], developed by the Foundations of Software Engineering group at Microsoft Research. AsmL is a strongly typed language based on the concepts of ASMs but also incorporates numerous object-oriented features and constructs for rapid prototyping of component-oriented software, thus departing in that respect from the theoretical model of ASMs; rather it comes with the richness of a fully fledged programming language. At the same time, AsmL lacks any built-in support for dealing with distributed systems.

Being deeply integrated with the software development, documentation, and runtime environments of Microsoft, its design was shaped by practical needs of dealing with fairly complex requirements and design specifications for the purpose of software testing; as such, it is oriented towards the world of code. This has made it less suitable for initial modeling at the early stages of design and also restricts the freedom of experimentation.

AsmL was first released as a stand-alone compiler that would compile AsmL specifications into Microsoft Windows executable files. The release also included plugin for Microsoft Word to support literate programming; i.e., one could include AsmL specifications into a Microsoft Word document using a special formatting style and then compile the Microsoft Word document using the AsmL compiler.

AsmL is now maintained as a community project at [102]. The experience of the design and applications of AsmL later contributed to the design of the Spec# programming language (an extension of the object oriented language C#) and the Spec Explorer software development tool [33, 103] which are both based on the Microsoft .NET framework. Spec Explorer is a tool for model-based specification and conformance testing of reactive, object-oriented software systems. Since 2004, AsmL has been integrated into Spec Explorer as one of the two languages supported by this tool. However, the official language of Spec Explorer is Spec# which has no resemblance to an ASM language.

3.6 Asmeta

The *Asmeta* project [60, 67], one of the most feature-rich ASM tool projects currently maintained, focuses on defining a metamodel for ASMs, called *Abstract State Machine Metamodel* or *AsmM* for short, based on the Model-Driven Engineering (MDE) [114] guidelines. In fact,

	ALMA	**AsmGofer**	**Xasm**	**ASM Workbench**	**AsmL**	**Asmeta**
Language	DASL	AsmGofer	Xasm	ASM-SL	AsmL	AsmetaL
ASM Support*	subset of basic	basic, dist., ND	basic, ND	basic, ND	basic, ND	basic, dist., ND
GUI	No	Yes	Yes	Yes	Yes, through .NET	Yes
Implementation	Prolog	Gofer	C source	ML and C++	.NET framework	Java
Extensibility	No	No	external C and Java functions	loose coupling through process communication	interaction with the .NET framework	external Java functions
Maintained	No	Yes	No	No	Yes	Yes
Available	No	Yes	Yes	No	Yes	Yes
Special Features	-	automatic and user-defined GUI	regular expressions, grammar definition, LaTeX support	pattern matching, model checking through SMV	integrated into the Spec Explorer modeling suite	various tools around an OMG metamodel of ASM

* *basic:* basic ASMs; *dist.:* distributed ASMs; *ND:* non-determinism

Table 3.1: Comparing ASM Tools and Languages

Asmeta is an instantiation of the OMG metamodeling framework [1] for ASMs. The main goals of the Asmeta project are to develop a unified abstract notation of ASMs independent of any specific implementation or concrete syntax and to develop a general framework for a wide interoperability and integration of ASM tools [67].

The Asmeta tool set is a collection of various components and tools around the *AsmM metamodel* that abstractly represents concepts and constructs of ASMs as described in [25]. The Asmeta language, called *AsmetaL*, is a textual notation for AsmM defined in terms of an EBNF grammar and serves as a fully typed language implementing all the constructs of basic, structured, and multi-agent ASMs as defined in [25]. AsmetaL is enriched by a *standard library* which is a collection of predefined ASM domains (such as Boolean, Integer, String, etc.) and functions and operations defined on those domains. A text-to-model *compiler* is also provided to parse and translate specifications written in AsmetaL language into AsmM instances (AsmM models) which can be executed by the *AsmetaS* simulator, implemented in Java. A graphical front-end and IDE for AsmetaL, called *Asmee*, is also provided in form of an Eclipse[2] plugin. The *ASM Tests Generation Tool* of Asmeta helps with creating test cases for AsmM models. Finally, a *graphical notation* is also provided as an alternative concrete syntax for AsmM.[3]

Since Asmeta is a metamodel-based framework, various tools and applications can be developed utilizing the already implemented features. AsmetaL is not an extensible language but external Java functions can be used to model static ASM functions.

3.7 Alternative Tools

In the first chapter of this thesis, we addressed alternative approaches towards modeling of computational systems. Abstract state machines, among many others, fall into the category of state-based formal methods that view the states of a system in terms of mathematical structures. In this category, one can point to methods such as Alloy, B, CASL, and the Z notation as four of the most popular approaches that share many similar concepts such as offering abstract notations, supporting declarative descriptions of system behavior in terms of constraints, and relying on tool support for analysis of specifications.

[2] http://www.eclipse.org

[3] As of writing this thesis, the Asmeta graphical editor is still under development.

The Common Algebraic Specification Language, or CASL, is a general purpose specification language based on first-order logic with induction. Different extensions of the language have been designed for specification and development of various kinds of systems such as reactive or concurrent [10]. The language is supported by a number of tools for checking the correctness of specifications, proving certain properties of models, and managing the formal software development process. Currently, The *Heterogeneous Tool Set*,[4] or *Hets* for short, seem to be the mainstream central toolset for CASL. It's a free software, with a license similar to the GNU General Public License [61], offering parsing, analysis, and prover integration for CASL and its extensions.

The Z notation [116] is a formal specification language designed with proofs in mind; it is based on axiomatic set theory, lambda calculus, and first-order predicate logic. There are quite a number of tools available for Z, most of them focused on theorem proving such as *ProofPower*[5], a suite of tools supporting specification and proof in the Z notation, *Z/Eves* [34], a front-end for the *Eves* theorem prover, and *HOL-Z*,[6] a proof environment for Z specifications based on the generic theorem prover Isabelle/HOL[7]. A free and open source animator for Z specifications, called *Jaza*,[8] is also available for evaluation, testing and (for some specifications) also execution of Z specifications.

Inspired by Z, the Alloy specification language [84] is designed as a light-weight formal specification language with the goal of providing fully automatic analysis of software specifications. However, unlike Z, Alloy's data structures are strictly first order. Alloy comes with *AlloyAnalyzer*, a model-checker that checks certain properties of specifications by exploring the states of the system and looking for execution instances that satisfy the properties (*simulation*) or by finding counterexamples that violate them (*checking*).

Among these approaches, the B method [2] has a more operational flavor and is the most similar approach to ASMs. It is essentially an abstract machine notation with a well-defined notion of refinement that facilitates transformation of abstract models into implementation. B comes with a rich set of both commercial and open source tools. Commercial tools such

[4] http://www.dfki.de/sks/hets

[5] Registered trademark of Lemma 1 Ltd., http://www.lemma-one.com/ProofPower

[6] http://www.brucker.ch/projects/hol-z

[7] http://www.cl.cam.ac.uk/research/hvg/Isabelle

[8] http://www.cs.waikato.ac.nz/~marku/jaza/

as *Atelier-B*[9] and the *B-Toolkit*[10] are available providing syntax analysis, theorem proving, and automatic refinement of B specifications down to implementation. A single-user free version of Atelier-B, called *B4Free*, is also available for the academic environment. There are also model checkers available for B; for example, *ProB*[11] offers fully automatic animation of many B specifications and can be used to systematically check a specification for errors.

[9] http://www.atelierb.eu

[10] http://www.b-core.com/btoolkit.html

[11] http://users.ecs.soton.ac.uk/mal/systems/prob.html

Part II

Design and Specification of CoreASM

Chapter 4

CoreASM: Architectural Overview

The CoreASM language and supporting tool architecture focus on early phases of the software design process. In particular, the goal is to encourage rapid prototyping with ASMs, starting with mathematically-oriented, abstract and untyped models and gradually refining them down to more concrete versions—a powerful technique for specification with refinement that has been exploited in [25] and [18]. In this process, we aim at maintaining executability of even fairly abstract models. Another important characteristic that differentiates our endeavor from previous experiences is the emphasis that we are placing on extensibility of the language. Historical developments have shown how the original, basic definition of ASMs from the Lipari Guide [74] has been extended many times by adding new rule forms (e.g., **choose**) or syntactic sugar (e.g., **case**). At the same time, many significant specifications need to introduce special backgrounds[1], often with non-standard operations. We want to preserve in our language the freedom of experimentation that has proven so fruitful in the development of ASM concepts, and, to this end, we have designed our architecture around the concept of *plugin*s that allows to customize the language to specific needs.

The architecture of the CoreASM engine is partitioned along two dimensions (see Figure 4.1).[2] The first one identifies the main components of the CoreASM engine and their relationships: a *parser*, an *interpreter*, a *scheduler*, and an *abstract storage* (Figure 4.2). We will discuss these components in more detail in Section 4.1. The second dimension, discussed in Section 4.3, distinguishes between what is in the *kernel* of the system—thus

[1] We call *background* a collection of related domains and relations packaged together as one logical unit.

[2] This chapter builds on and significantly extends what we have previously published in [48, Section 2].

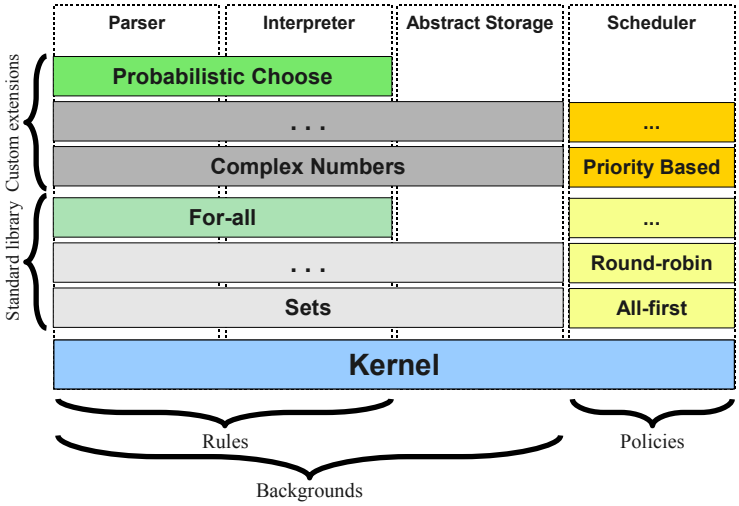

Figure 4.1: Layers and Modules of the CoreASM Engine

implicitly defining the extreme bare bones of the model—and what is instead provided by extension plugins.

These two dimensions correspond to what in the ASM literature have been called *modular decomposition* and *conservative refinement* respectively [18].[3] In particular, our plugins progressively extend (potentially in a conservative way) the capabilities of the language accepted by the CoreASM engine, in the same spirit in which successive layers of the Java [118] and C# [19] languages have been used to structure the language definition into manageable parts.

In this chapter we provide an overview of the architecture of the CoreASM engine and present its components. We also explore the execution lifecycle of the engine and its control state model, and discuss the micro-kernel approach to the design of the engine and its extensibility mechanisms.

4.1 CoreASM Components

The CoreASM engine consists of four components: a parser, an interpreter, a scheduler, and an abstract storage (Figure 4.2). The interpreter, the scheduler, and the abstract storage

[3] While CoreASM plugins are expected to extend the engine mostly through a conservative refinement, the CoreASM architecture does not restrict the plugins to such a refinement.

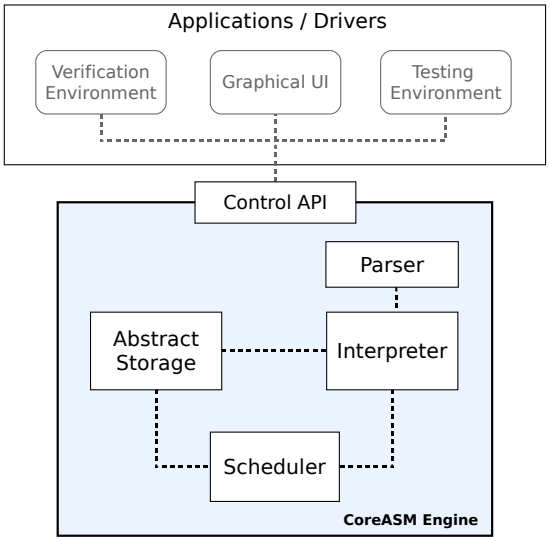

Figure 4.2: Overall Architecture of CoreASM

work together to simulate an ASM run. The engine interacts with the environment through a single interface, called the *Control API*, which provides various operations such as loading a CoreASM specification, starting an ASM run, or performing a single step.

The parser reads a CoreASM specification and generates annotated abstract syntax trees for rules (programs) and definitions of the specification. Each node in these trees may have a reference to the plugin that provides the corresponding syntax. For example, in Figure 4.3, there are nodes that belong to the backgrounds of sets and Booleans; this information will be used by the interpreter and the abstract storage to perform operations on these nodes with respect to the background each node comes from.

The interpreter, executes programs and rules, possibly calling upon background plugins to perform expression evaluation, and upon rules plugins to interpret certain rule forms. It obtains an annotated parse tree from the parser and generates a multiset of *update instructions*, each of which represents either an update, or an arbitrary instruction which will be processed at a later stage by corresponding plugins to generate actual updates (as will be described in more detail on page 59)[4]. The interpreter interacts with the abstract storage to retrieve data from the current state and by executing statements it gradually

[4]Where no confusion can arise, in the rest of this thesis we use the generic term "updates" to refer both to actual updates and to update instructions.

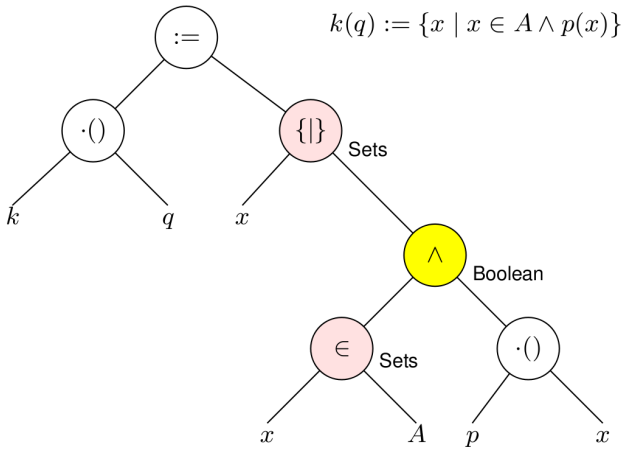

Figure 4.3: Sample Annotated Parse Tree

creates the update set leading to the next state.

The abstract storage manages the data model for the abstract state; in particular, it maintains a representation of the current state of the machine that is being simulated. The state is modeled as a map from locations to opaque elements of a universe ELEMENT. The abstract storage also provides interfaces to retrieve values from a given location in the current state and to apply updates. To evaluate a program, the interpreter interacts with the abstract storage in order to obtain values from the current state and generates updates for the next state. In addition, abstract storage also provides auxiliary information about the locations of the current state, such as the ranges and domains of functions or the background to which a particular function or value belongs to.

Finally, the scheduler orchestrates every computation step of an ASM run. In a basic ASM, the scheduler merely arranges the execution of a step: it receives a *step* command from the Control API, invokes the interpreter, and instructs the abstract storage to aggregate the update instructions and *fire* (apply to the state) the resulting update set (if consistent) when the interpreter finishes the evaluation of the program. It then notifies the environment through the Control API of the results of the step.

For distributed ASMs [25], the scheduler also organizes the execution of agents in each computation step. At the beginning of each DASM computation step, the scheduler chooses a subset of agents which will contribute to the next computation step of the machine. The scheduler directly interacts with the abstract storage to retrieve the current set of agents, to assign the current executing agent, and to collect the update set generated by

the interpretation of all the agents' programs. Updates are then fired and the environment is notified as for the previous case.

4.2 Engine Lifecycle

The process of executing a CoreASM specification in the CoreASM engine consists of the following steps:

1. Initializing the engine (Figure 4.4)
   (a) Initializing the kernel
   (b) Loading the plugins library catalogue
   (c) Loading and activating core plugins
2. Loading a CoreASM specification (Figure 4.5)
   (a) Parsing the specification header
   (b) Loading required plugins as declared in the specification
   (c) Parsing the specification body
   (d) Initializing the abstract storage
   (e) Setting up the initial state[5]
3. Execution of the specification
   (a) Execute a single step
   (b) If termination condition is not met, repeat from 3a.

The execution process of a single step in the CoreASM engine is as follows (refer also to Figures 4.6 to 4.9 in Section 4.2): The Control API sends a *step* command to the scheduler. *(i)* The scheduler gets the whole set of agents from the abstract storage. *(ii)* It selects a subset of these agents to participate in the next computation step. *(iii)* One by one, the scheduler selects and removes agents from this set and assigns them to the special variable *self* in the abstract storage.[6] *(iv)* The scheduler then calls the interpreter to run the program of the current agent (retrieved by accessing *program(self)* in the current state). *(v)* The

[5] This ensures that there is at least one agent in the state, the program of that agent being the rule marked with **init** and that agent will contribute to the first step of the simulation.

[6] This is done implicitly by assigning the agent as the value of *executingAgent*. See Section 4.2.3.

interpreter evaluates the program.[7] *(vi)* When the evaluation of the program is complete, the interpreter notifies the scheduler. *(vii)* The scheduler gathers the computed update set and repeats from step (iii) until there is no agent left in the set. When all the agents are executed, the scheduler calls the abstract storage to apply the accumulated updates to the state. *(viii)* If the update set is inconsistent, the abstract storage notifies the scheduler and the notification may lead to selection of a different subset of agents to be executed.[8] If the update set is applied successfully, the Control API is notified of the successful step.

At the end of the execution of each step, the resulting state is optionally made available by the abstract storage module for inspection through the Control API. The termination condition can be set through the user interface of the CoreASM engine, choosing between a number of possibilities (e.g., a given number of steps are executed; no updates are generated; the state does not change after a step; an interrupt signal is sent through the user interface).

In the following sections, we present a high-level but precise specification of the execution process which was presented informally at the beginning of this section. The structure of the specification is that of a control state ASM [25, Sec. 2.2.6][9], as shown in Figures 4.4 to 4.9. The current state of such ASM is given by the variable *engineMode* that controls the execution of rules at any step. The ASM rules corresponding to the control state ASM are also presented.

4.2.1 Engine Initialization

The CoreASM engine starts its execution in the *Idle* state (Figure 4.4). In this state, the engine simply waits for a control command, such as *init* or *step*, from the environment which could be an interactive GUI or a debugger, to start the corresponding task.

Receiving an *init* command (Figure 4.4) will change the state of the engine to *Initializing Kernel* in which the engine initializes its kernel, loads its plugin catalog (the set of all the plugins available to the engine), and finally loads the core plugins. The following rules in Control API abstractly define these tasks. We refer the reader to Section 5.5 for more details on loading plugins.

---

[7] This may include a series of interactions between the interpreter and the abstract storage to get values from the current state, which in turn may require interpreting other code fragments, e.g., for derived functions.

[8] The engine can also report (e.g. in a log file) the set of agents whose updates produced an inconsistent update set.

[9] In fact we are using a variant of control state ASMs; see Section 5.5.5 for more details.

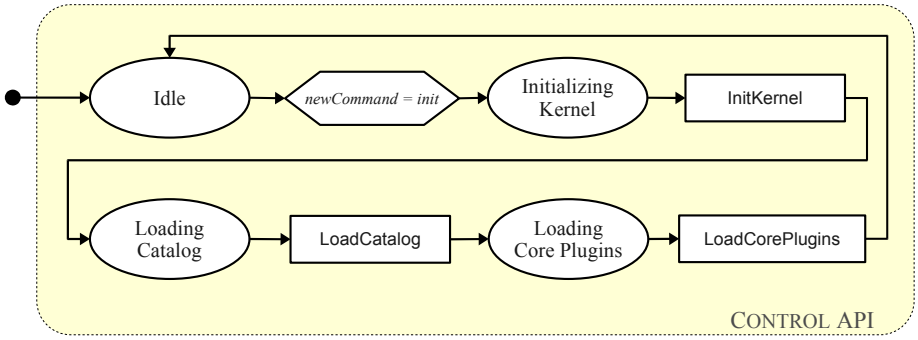

Figure 4.4: Control State ASM of Initializing CoreASM Engine

Control API

InitKernel ≡
  pluginCatalog := {}
  loadedPlugins := {}
  grammarRules := {}
  specification := *undef*
  isStateInitialized = *false*

LoadCatalog ≡
  forall *pName* **in** *availablePlugins* **do**
    let *p* = *createPlugin*(*pName*) **in**
      add *p* **to** *pluginCatalog*

LoadCorePlugins ≡
  forall *p* **in** *corePlugins* **do**
    LoadPlugin(*p*)

In order to keep the model consistent, some of the functionalities of the CoreASM kernel can be encapsulated in special *core plugins*. For example, in Section 5.3 we will see how plugins can contribute to the aggregation of updates after every computation step. However, there is also a default aggregation behavior that must be provided by the kernel itself. By encapsulating that default behavior in a special core plugin (*Kernel plugin*), we are able to reduce the complexity of the aggregation process and specify it in a simple and concise form. So far, the set *corePlugins* consists of only one plugin; i.e. *corePlugins* = {*kernelPlugin*}.

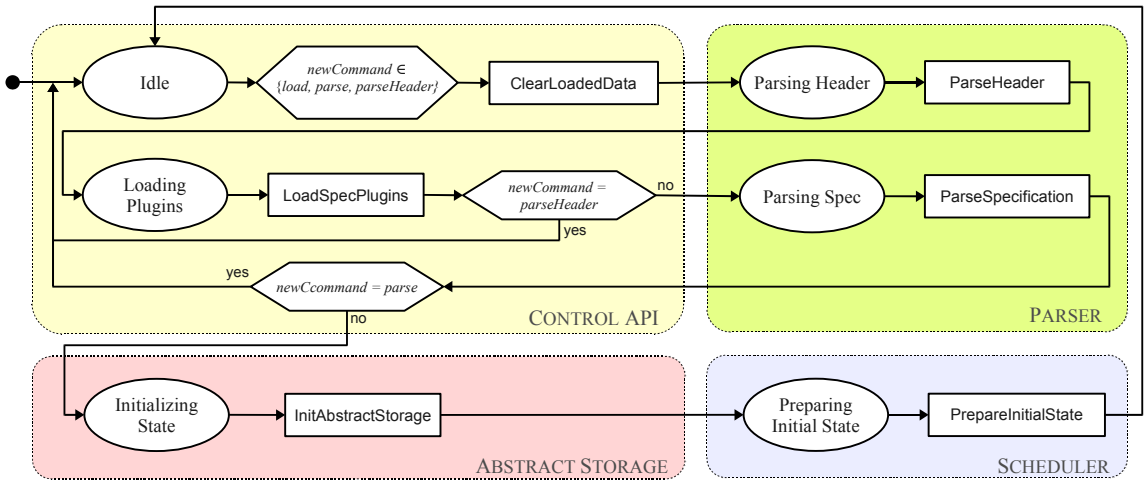

Figure 4.5: Control State ASM of Loading a CoreASM Specification

4.2.2 Loading Specification

Receiving a *load* command causes the engine to load a new specification (Figure 4.5). The engine first clears previously loaded data, reads the specification file and then parses the specification header to get the list of specific plugins required to be loaded.

Control API

ClearLoadedData ≡
  if *specHasBeenLoaded* **then**
    seq
      loadedPlugins := {}
      grammarRules := {}
      specification := *getSpecification*(*newCommand*)
    next
      LoadCorePlugins
where
  specHasBeenLoaded ≡ |*loadedPlugins*| > |*corePlugins*|

Parser

ParseHeader ≡
  specPlugins := *requestedPlugins*(*specification*)

Loading the required plugins is done in two steps. First, all the package plugins (plugins

that are basically a set of other plugins) are expanded and their enclosed plugins are added to the list of required plugins. In the next step, plugins are loaded one by one according to their loading priority.

Control API

LoadSpecPlugins ≡

```
seq
// 1. expanding package plugins
  forall p in specPlugins do
    if isPackagePlugin(p) then
      forall p' in enclosedPlugins(p) do
        add p' to specPlugins
next
// 2. loading plugins with the maximum load priority first
  while |specPlugins\loadedPlugins| > 0 do
    let toLoad = specPlugins\loadedPlugins in
      choose p in toLoad with maxPriority(p, toLoad) do
        if requiredPlugins(p) ⊂ specPlugins then
          LoadPlugin(p)
        else
          Error('Cannot load plugin.')
```

After all the required plugins are loaded, the specification is parsed using the grammar rules provided by the plugins. The root node of the resulting parse tree is kept for future references.

Parser

ParseSpecification ≡

rootNode(*specification*) ← Parse(*specification*, *grammarRules*)

To prepare the engine for the first simulation step, Abstract Storage is initialized taking into account all plugins contributions, such as backgrounds, universes, functions, and macro rules. A universe of *Agents* and a function *program* that assigns programs to agents are also created in this step. See page 92 for the definition of LoadVocabularyPlugins.

Abstract Storage

```
InitAbstractStorage ≡
  let newState = new(STATE) in
    state := newState
    InitializeState(newState)
    LoadVocabularyPlugins(newState)

InitializeState(state) ≡
  let u = new(UNIVERSEELEMENT) in
    stateUniverse(state, "Agents") := u
  let f = new(FUNCTIONELEMENT) in
    stateFunction(state, "program") := f
  executingAgent := undef      // holds the value of 'self' in the simulated machine
  stepCount := 0
```

Finally, an initial state is created with at least one agent that, in the first step of the simulation, will run the **init** rule as its main program. In addition, based on the set of plugins used by the specification, a scheduling policy will also be chosen by the scheduler.[10]

Scheduler

```
PrepareInitialState ≡
  LoadSchedulingPolicy
  let a = new(ELEMENT) in
    initAgent := a
    SetValue(("Agents", ⟨a⟩), true_e)
    SetValue(("program", ⟨a⟩), initRule)
```

Alternatively, an external application may ask the engine to only *parse* the specification (and not loading it). This is useful when an application needs to use only the parsing functionality of the engine, for example to work on a parse-tree view of a specification. In this case, the last two steps of initializing state and preparing the initial state will be skipped. Also, an application can query the list of plugins required by a given specification by sending a *parseHeader* command. In this case, the engine does not parse the specification and stops after loading the required plugins.

[10] We refer the reader to Appendix A.3 for more details.

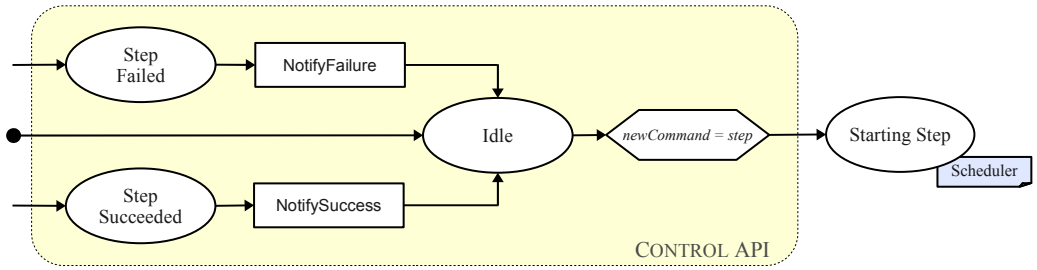

Figure 4.6: Control State ASM of a *step* command: Control API Module

4.2.3 Execution of Specification

A *step* command triggers the start of a computation step; this is performed by changing the control state to *Starting Step* which then transfers the control flow to the scheduler.

The StartStep rule in the scheduler initializes *updateInstructions* (the multiset of accumulated update instructions for the current step) and *selectedAgentsSet* (the set of agents selected to perform computation in the current step) and assigns the current set of agents in the simulated machine to *agentSet* by querying the abstract storage module for the current value of *Agents* and only picking those agents whose program is not undefined. We model the query process through the abstract function $getValue(l)$ which takes a location l and retrieves the value of the location from the simulated state. We use the notation "term" to denote the quoted variable or literal term *term* in the simulated machine. Based on the retrieved set of agents, a new schedule is then created by CreateSchedule. The control state is then changed to *Selecting Agents*.

Scheduler

StartStep $\equiv$

$updateInstructions := \{\!\{\}\!\}$

$selectedAgentsSet := \{\}$

if $stepCount < 1$ **then**

$agentSet := \{initAgent\}$

else

$agentSet := \{a | a \in getValue((\text{“Agents”}, \langle\rangle)) \wedge getValue((\text{“program”}, \langle a \rangle)) \neq \text{undef}_e\}$

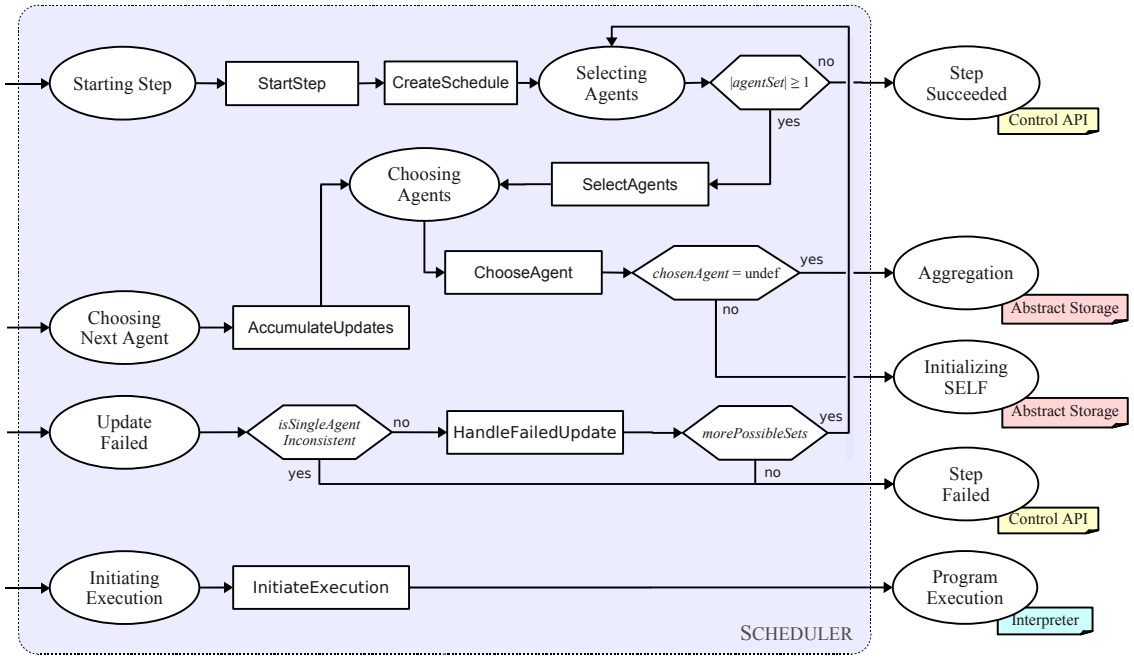

Figure 4.7: Control State ASM of a *step* command : Scheduler

CreateSchedule ≡
  if *schedulingPolicy* ≠ *undef* **then**
    let $R = newScheduleRule(schedulingPolicy)$ **in**
      $schedule \leftarrow R(schedulingGroup, agentSet)$

In the *Selecting Agents* state, if no agent is available to perform computation, the step is considered complete; otherwise, the SelectAgents rule chooses a set of agents to execute in the current step. If there is no scheduling policy provided by any of the plugins, a non-deterministic subset of the agents is chosen; otherwise, the selected agents will be determined by the current scheduling policy. The ChooseAgent rule chooses an agent from this set and changes the state to *Initializing SELF* which leads to the execution of the SetChosenAgent rule in the abstract storage module. After the execution of the agent, the computed updates are accumulated by AccumulateUpdates rule in the *Choosing Next Agent* state, and control state is changed back to *Choosing Agent* until all selected agents have been executed.

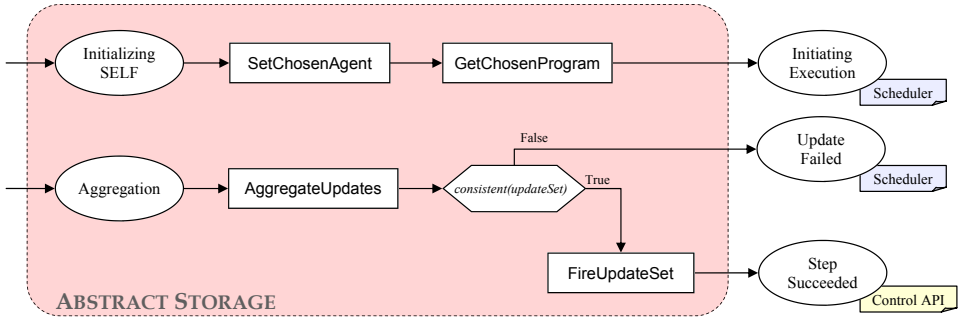

Figure 4.8: Control State ASM of a *step* command : Abstract Storage

---

Scheduler

SelectAgents ≡
  if *schedulingPolicy* = *undef* **then**
    choose *s* **with** $s \subseteq agentSet \wedge |s| \geq 1$ **do**
      selectedAgentsSet := *s*
  else
    selectedAgentsSet := *head*(*schedule*)
    schedule := *tail*(*schedule*)

ChooseAgent ≡
  choose *a* **in** *selectedAgentsSet* **do**
    remove *a* **from** *selectedAgentsSet*
    chosenAgent := *a*
  ifnone
    chosenAgent := *undef*

AccumulateUpdates ≡
  add *updates*(*root*(*chosenProgram*)) **to** *updateInstructions*

Two rules in the abstract storage module take care of setting the chosen agent and of retrieving the program associated with the chosen agent (by accessing *program*(*self*) in the simulated state). Control then moves back to the scheduler at *Initiating Execution*.

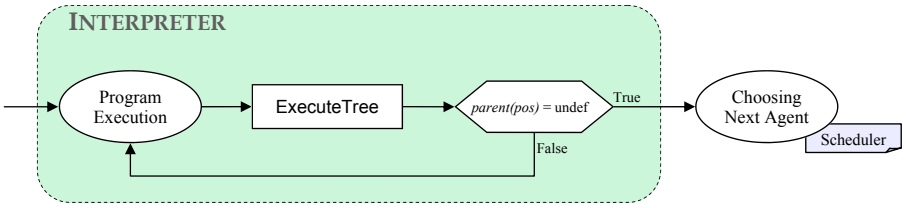

Figure 4.9: Control State ASM of a *step* command : Interpreter

---

Abstract Storage

SetChosenAgent ≡
  $executingAgent := chosenAgent$

GetChosenProgram ≡
  $chosenProgram := getValue((\text{“program”}, \langle executingAgent \rangle))$

The execution of the program of the chosen agent is initiated in the *Initiating Execution* state in the scheduler and then starts in the *Program Execution* state in the interpreter. During the execution, computed update instructions are progressively added to *updateInstructions*, and when all selected agents have performed their computation, control moves to *Aggregation* state in the abstract storage, where the final update set is calculated and then applied to the current state.

Extending the basic idea presented in [118], we interpret a program by associating values, updates and locations to nodes in the parse tree of the program. Before actually starting the interpreter, previously computed values are removed by the InitiateExecution rule, and the current position in the tree (denoted by the nullary function *pos*) is initialized to the root node of the tree that represents the current program (that is, the program of the current agent, as established above).

Scheduler

InitiateExecution ≡
  let $p = root(chosenProgram)$ **in**
    ClearTree(p)
    $pos := p$

The specification of the interpreter is explored in detail in Section 5.2. We do not include here the full specification for the interpreter; we show instead its most interesting feature,

that is the way it interacts with rule and background plugins to delegate interpretation of the associated extensions. To do this, we slightly extend the ASM framework to include ASM rules (programs) as elements of the state; i.e. we assume that ASM rules are elements of the domain RULE and that they can be treated as terms and so can be assigned as values of functions.

As already discussed earlier, nodes of the parse tree corresponding to grammar rules provided by a plugin are annotated with the plugin's identifier. The annotation process is done during parsing, but here we abstract from the details of how it is implemented, and use instead an oracle function *plugin*(*node*) for this purpose. While interpreting the parse tree (see ExecuteTree below), if a node is found to refer to a plugin, rules provided by that plugin are obtained through the *pluginRule* function and executed; otherwise, the kernel interpreter rules (see Section 5.2) are used. Results of the interpretation of node *pos* are stored alongside the node, and accessed by three functions: *value*(*pos*) returns the computed value for an expression node, *updates*(*pos*) returns the set of updates generated by a rule node, and *loc*(*pos*) returns the location denoted by the node (which is used as lhs-value for assignments). Section 5.2.1 presents a more precise definition of these functions.

Interpreter

ExecuteTree ≡
  if ¬*evaluated*(*pos*) **then**
    if *plugin*(*pos*) ≠ *undef* **then**
      let *R* = *pluginRule*(*plugin*(*pos*)) **in**
        R
    else
      KernelInterpreter
  else
    if *parent*(*pos*) ≠ *undef* **then**
      pos := *parent*(*pos*)

After executing the programs of all the selected agents, all the update instructions will have been accumulated in *updateInstructions*. Control will move from *Choosing Agent* in the scheduler to *Aggregation* in the abstract storage module. In the *Aggregation* state, the abstract storage aggregates update instructions to compute updates on the locations of the state (see Section 5.3.2 for details), checks the consistency of the computed updates (possibly interacting with the relevant background plugins to evaluate equality), and either applies

the updates to the current state through FireUpdateSet (thus obtaining the next state), or provides an indication of failure by changing the state to *Update Failed*.

Abstract Storage

AggregateUpdates ≡
  updateSet ← Aggregate(*updateInstructions*)

FireUpdateSet ≡
  forall $(l, v) \in$ *updateSet* **do**
    SetValue(l, v)

In the earlier versions of CoreASM [47], if an inconsistent set of updates would be generated in a step, the HandleFailedUpdate rule in the scheduler module would prepare a different subset of agents for execution, and the step would be re-initiated. As a result, if a single agent would produce inconsistent updates, instead of reporting the inconsistency as an error, that agent would be removed from the set of computing agents. We later improved the control flow so that an update fails if the inconsistent set of updates are produced by a single agent. Otherwise, if the inconsistency is between two updates from two different agents, other combinations of agents are tried and the process is iterated until either a consistent set of updates is generated, in which case the computation proceeds to the *Step Succeeded* state of the Control API, or all possible combinations have been exhausted, in which case controls moves to the *Step Failed* state. It should be noted that the selection will also consider subsets containing a single agent, so the process fails only when no agent can successfully perform a step.

Depending on the outcome of the previous stage, either of the rules NotifySuccess or NotifyFailure of the Control API notify the environment of the success or failure of the step, and return to the *Idle* state awaiting further commands from the environment (e.g., another *step* command to continue the computation).

Control API

NotifySuccess ≡
  stepCount := *stepCount* + 1

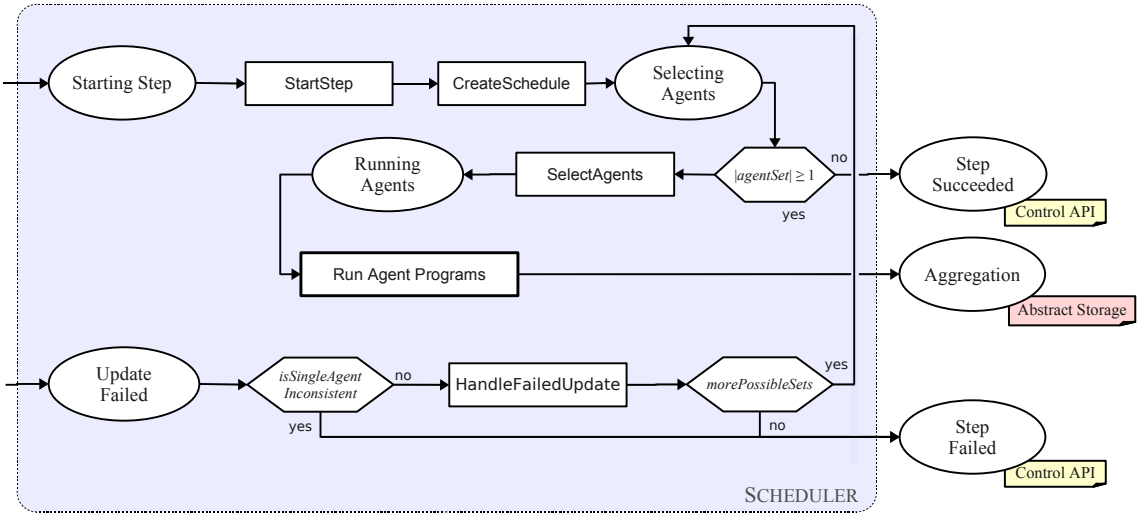

Figure 4.10: Revised Control State ASM of a *step* command: Concurrent Scheduler

4.2.4 Concurrently Running Agents

We can abstract away from the details of interleaved execution of selected agents in every step of the simulation and model the process in a parallel form. This abstraction is beneficial as it removes the unnecessary sequential order of the execution of agents, hence avoiding over-specification of the engine, and it allows for a more efficient implementation of the engine by *a)* removing the explicit control flow loop around the interpretation of single parse tree nodes (see Figure 4.9) and *b)* enabling concurrent execution of agents on multi-processor machines.

In order to run agent programs in parallel, every function and rule related to the interpretation of the programs should be parameterized by the agents accessing them. As a result, the control state diagram of the scheduler will be reduced to that of Figure 4.10. The RunAgentPrograms rule in the diagram will directly use a parameterized version of the ExecuteTree rule, thereby eliminating the control state diagram of the interpreter.

Scheduler

```
RunAgentPrograms ≡
  forall a ∈ selectedAgentsSet do
    let p = getValue((“program”, ⟨a⟩)) in
      seq
        pos(a) := root(p)
        ClearTree(p)
      seq
        while ¬isEvaluated(root(p)) do
          ExecuteTree(a)
      next
        add updates(root(p)) to updateInstructions
```

4.3 CoreASM Plugins

In keeping with the micro-kernel spirit of CoreASM, most of the functionality of the engine is implemented through plugins to a minimal kernel. In principle, there are three basic dimensions being considered for extending and altering CoreASM by means of plugins, respectively related to: (i) data structures, (ii) control structures, and (iii) the execution model.

i) The possibility of conveniently extending data structures as needed is extensively discussed in the theoretical ASM literature, e.g. in [13, 12], where the concept of *background* refers to an implicitly given part of an abstract machine state, assuming that it provides whatever standard means are normally supposed to be available in a given application context [13]. Plugins extending the data structures of the engine provide all that is needed to define and work with new backgrounds, namely (a) an extension to the parser defining the concrete syntax (operators, literals, static functions, etc.) needed for working with elements of the background; (b) an extension to the abstract storage providing encoding and decoding functions for representing elements of the background for storage purposes, and (c) an extension to the interpreter providing the semantics for all the operations defined in the background. The *Set* plugin, presented in Section 6.3.2, is an example of a background plugin (see Figure 4.1).

ii) Plugins can extend the control structures of CoreASM with respect to both new syntactic constructs that are semantically meaningful and those that only provide syntactic

sugar (i.e., the semantics of which could also be expressed by means of in-language transformations). These plugins provide specific rule forms, with the understanding that the execution of a rule always results in a (possibly empty) set of updates. Thus, they include (a) an extension to the parser defining the concrete syntax of the rule form; (b) an extension to the interpreter defining the semantics of the rule form.

iii) Finally, the need for altering or extending the execution model is justified by pragmatic considerations. The execution model refers to dynamic features of CoreASM, including scheduling policies, exception handling, and instrumentation of program execution for analytical purposes. Plugins can alter the execution model of the engine either by providing new scheduling policies to the scheduler, used to determine at each step the next set of agents to execute, or by extending the control sate ASM of the engine. See Section 5.5.5 for more details.

In CoreASM, the kernel (see Figure 4.1) only contains the bare essentials, that is, all that is needed to execute only the most basic ASM. As the state of an ASM machine is defined by functions and universes, the two domains of *functions* and *universes* are included in the kernel. Universes are represented through their characteristic functions, hence *Booleans* are also included in the kernel. As an ASM program is defined by a finite number of rules, the domain of *rules* is also included in the kernel. It should be noted that the kernel includes the above mentioned domains, but not all of the expected corresponding backgrounds. For example, while the domain of Booleans (that is, true and false) is in the kernel, the Boolean algebra ($\wedge$, $\vee$, $\neg$, etc.) is not, and is instead provided through a background plugin. In the same vein, while universes are represented in the kernel through set characteristic functions, the background of finite sets is implemented in a plugin, which provides expression syntax for defining them (see the example in Figure 4.3), as well as an implicit representation for storing sets in the abstract state, and implementations of the various set theoretic operations (e.g., $\in$) that work on such implicit representation.

The kernel includes only two types of rules: assignment and **import**. This particular choice is motivated by the fact that without updates established by assignments there would be no way of specifying how the state should evolve, and that **import** has a special role in introducing new elements to the state. All other rule forms (e.g., **if**, **choose**, **forall**), as well as sub-machine calls and macros, are implemented as plugins in a standard library.

Finally, there is a single scheduling policy implemented in the kernel, namely the pseudo-random selection of an arbitrary set of agents at a time, which is sufficient for multi-agent ASMs where no assumptions are made on the scheduling policy.

As already mentioned, the CoreASM engine is accompanied by a *standard library* of plugins including the most common backgrounds and rule forms (i.e., those defined in [25]), an extension library including a small number of specialized backgrounds and rules, and by a set of specifications for writing new plugins that can easily be integrated in the environment. Extension plugins must be explicitly imported into an ASM specification by an explicit **use** directive.

The plugin framework is further discussed in Section 5.5.

Chapter 5

CoreASM: The Kernel

In this chapter, we look into the details of the CoreASM kernel and its four components. We formally define the interfaces of these components in form of functions and operations (ASM rules). In case of the Abstract Storage, we present the initial structure of simulated *states* in CoreASM and formally define the elements of which it consists of. We then provide a detailed specification of the Interpreter, building on the ExecuteTree rule we presented in Section 4.2. In Section 5.3, we look into the concepts of rules and updates in CoreASM and finally conclude this chapter with an overview of the CoreASM plugin framework.

5.1 The Abstract Storage

Abstract Storage maintains a representation of the current state of the simulated machine in CoreASM. In order to distinguish between the values in the simulated state and the values in our ASM model of the engine, we denote the values of the simulated state as *elements* modeled by the domain ELEMENT. There is a special element in the state that represents the *undefined* value or *undef*. Henceforth, this element is denoted by undef_e .

Elements can belong to different backgrounds, such as Set, Number, Map, and so on. The background of every element is defined by the following function whose default value is "Element" for all elements that do not belong to a particular background:

$$bkg : \text{ELEMENT} \mapsto \text{NAME}$$

The kernel also defines a notion of equality on elements which can be extended by plugins providing special backgrounds. For any two elements e_1 and e_2, the notion of equality is

defined as:[1]

$$equal(e_1, e_2) \equiv equal_{bkg(e_1)}(e_1, e_2) \vee equal_{bkg(e_2)}(e_2, e_1)$$

providing that[2]

$$\forall e_1, e_2 \in \textsc{Element} \quad equal_{Element}(e_1, e_2) \equiv e_1 = e_2$$

We model the simulated abstract state as an element of the domain STATE where every $s \in \textsc{State}$ in principle models a mapping from locations to values (elements). We have:

$$content : \textsc{State} \times \textsc{Location} \mapsto \textsc{Element}$$

During a simulation, the current simulated state is represented by the nullary function *state*: STATE. Locations are values of the domain LOCATION and each represents a pair of function name and a sequence of arguments:

$$name_{lc} : \textsc{Location} \mapsto \textsc{Name}$$
$$args_{lc} : \textsc{Location} \mapsto \textsc{List}(\textsc{Element})$$

We often denote locations by a pair $(f, \langle a_1, \ldots, a_n \rangle)$ where f is the name of the location and $\langle a_1, \ldots, a_n \rangle$ are the arguments.

In addition to its *content*, a CoreASM state also consists of backgrounds, universes, functions and rules. Before we look into functions and universes, we introduce Boolean elements, the most basic type of elements in the state.

Boolean Elements

We model Boolean elements by values of the domain BOOLEANELEMENT which has only two elements true_e and false_e, respectively representing Boolean values *true* and *false*. The following functions map Boolean elements to Boolean values and vice versa.

$$booleanElement : \textsc{Boolean} \mapsto \textsc{BooleanElement}$$
$$booleanValue : \textsc{BooleanElement} \mapsto \textsc{Boolean}$$

[1] Here, the notation $f_x(a_1, \ldots, a_n)$ can be seen as a syntactic sugar for $f(x, a_1, \ldots, a_n)$ and if x is missing, it can be interpreted as $f(undef, a_1, \ldots, a_n)$.

[2] In this equation, *Element* refers to the background name "Element".

For example, we have:

$$booleanElement(true) = \mathsf{true}_e$$
$$booleanValue(\mathsf{true}_e) = true$$

Equality of Boolean elements are simply defined based on the equality of the Boolean values they represent:

$$equal_{Boolean}(b_1, b_2) \equiv booleanValue(b_1) = booleanValue(b_2)$$

For all $b \in$ BooleanElement we have $bkg(b) =$ "Boolean".

Function Elements

Functions defined in a CoreASM state are modeled by function elements, values of the domain FunctionElement. Every CoreASM state holds a mapping of function names to function elements:

$$stateFunction : \text{State} \times \text{Name} \mapsto \text{FunctionElement}$$
$$functions : \text{State} \mapsto \text{Set}(\text{FunctionElement})$$
$$functions(s) \equiv \{f \mid f \in \text{FunctionElement} \wedge (\exists n \in \text{Name}, stateFunction(s, n) = f)\}$$

Function elements in principle represent a mapping from a sequence of elements (arguments of the function) to an element (the value of the function for those arguments):

$$value_{fe} : \text{FunctionElement} \times \text{List}(\text{Element}) \mapsto \text{Element}$$

ASM functions are classified into six categories of *monitored* (or *in*), *controlled*, *shared*, *out*, *static*, and *derived*. Monitored functions, or input functions, are those whose values are only read but never updated by the machine and can only be updated by the environment. Controlled functions, are the opposite; their values can be updated only by the machine and not the environment. Shared functions can be updated and read by both the machine and the environment. The values of out functions can only be updated but never read by the machine; they are intended for output and their values can be read by the environment of the machine. Static functions are constants and their values never change in course of an ASM run. Derived functions can be read by both the machine and the environment, but cannot be updated; their values are defined by a fixed scheme in terms of other functions. In CoreASM, classes of function elements are defined by the following function whose default

value is *controlled*:[3]

$$class_{fe} : \text{FUNCTIONELEMENT} \mapsto \{monitored, controlled, out, static, derived\}$$

Hence, modifiability of a function element f is defined as follows:

$$isModifiable(f) \equiv class_{fe}(f) \in \{controlled, out\}$$

If a function element is modifiable, its value for a particular sequence of arguments can be assigned by the following rule:

Abstract Storage

SetValue$_{\text{fe}}$ $(f, args, v) \equiv$
  if $isModifiable(f)$ **then**
    $value_{fe}(f, args) := v$

Every function element f is also a member of ELEMENT and $bkg(f) =$ "Function". Finally, two function elements are considered to be equal, if for all the possible arguments, they hold the same values.[4] For all $f_1, f_2 \in$ FUNCTIONELEMENT, we have:[5]

$$equal_{Function}(f_1, f_2) \equiv \forall a \in \text{LIST(ELEMENT)} \quad value_{fe}(f_1, a) = value_{fe}(f_2, a)$$

To retrieve the value of a function, the following derived function is defined as part of the interface of Abstract Storage:

$$getValue : \text{LOCATION} \mapsto \text{ELEMENT}$$
$$getValue(l) = \begin{cases} value_{fe}(\mathcal{F}, args_{lc}(l)), & \text{if } value_{fe}(\mathcal{F}, args_{lc}(l)) \neq undef; \\ \text{undef}_e, & \text{otherwise.} \end{cases}$$

where $\mathcal{F} = stateFunction(state, name_{lc}(l))$. The *getValue* function is later refined in Appendix A.1. In addition, Abstract Storage also provides the following macro rule to set the

[3] CoreASM does not support *shared* functions at this point.

[4] Since this definition is not necessarily computable, in practice we assume any two distinct functions to be unequal, unless defined otherwise (e.g., see Section 6.3.7). Hence, we have:

$$\forall f_1, f_2 \in \text{FUNCTIONELEMENT} \quad equal_{Function}(f_1, f_2) \equiv f_1 = f_2$$

[5] In ASMs, all functions are total. Partial functions are turned into total functions by introducing a xed special value *undef* and interpreting $f(x) = undef$ as $f(x)$ being undened. [25]

value of a location in the state:

Abstract Storage

SetValue$(l, v) \equiv$
  let $\mathcal{F} = stateFunction(state, name_{lc}(l))$ **in**
    if $\mathcal{F} \neq undef$ **then**
      $\mathsf{SetValue}_{\mathsf{fe}}(\mathcal{F}, args_{lc}(l), v)$

Universe Element

Universe elements, values of domain UNIVERSEELEMENT, represents the universes defined in a CoreASM state. Every CoreASM state holds a mapping of universe names to universe elements defined in that state:

$$stateUniverse : \text{STATE} \times \text{NAME} \mapsto \text{UNIVERSEELEMENT}$$
$$universes : \text{STATE} \mapsto \text{SET}(\text{UNIVERSEELEMENT})$$
$$universes(s) \equiv \{u \mid u \in \text{UNIVERSEELEMENT} \wedge (\exists n \in \text{NAME}, stateUniverse(s, n) = u)\}$$

Since universes are sets of elements (or values in ASM), we model them by their set characteristic functions. Hence, every universe element is also a function element. We have:

$$\forall u \in \text{UNIVERSEELEMENT},\ u \in \text{FUNCTIONELEMENT}$$

To conveniently view universe elements as sets, we define a membership function on universes:

$$member_{ue} : \text{UNIVERSEELEMENT} \times \text{ELEMENT} \mapsto \text{BOOLEAN}$$

For example, if element e belongs to the universe u in the current state of the simulated machine, we have $member_{ue}(u, e) = true$. As a result, for every $u \in$ UNIVERSEELEMENT and every $e \in$ ELEMENT, we have

$$value_{fe}(u, e) \equiv booleanElement(member_{ue}(u, e))$$
$$\mathsf{SetValue}_{\mathsf{fe}}(u, \langle e \rangle, b) \equiv member_{ue}(u, e) := booleanValue(b)$$

Equality of universes is defined as the equality of their characteristic functions:

$$\forall u_1, u_2 \in \text{UNIVERSEELEMENT}\ \ equal_{Universe}(u_1, u_2) \equiv equal_{Function}(u_1, u_2)$$

For all $u \in$ UNIVERSEELEMENT we have $bkg(u) =$ "Universe".

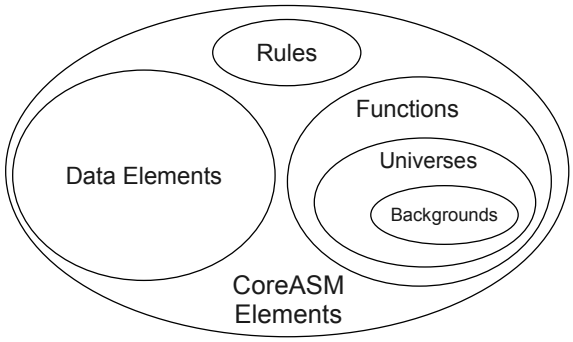

Figure 5.1: CoreASM Elements in the Kernel

Background Elements

In CoreASM, backgrounds are special universes with a static membership function. The assumption is that backgrounds contain all the elements they represent; e.g., background of sets represent all the possible sets. In principle, backgrounds represent "types" of elements mostly with internal structures. See, for example, how we define the backgrounds of character strings and sets in sections 6.2.3 and 6.3.2.

We model backgrounds by elements of the domain BACKGROUNDELEMENT. For every background element b, $newValue(b)$ must be defined to return a default element of that background; e.g., an empty set, an empty list, and such. We have:

$$newValue : \text{BACKGROUNDELEMENT} \mapsto \text{ELEMENT}$$
$$\forall\ b \in \text{BACKGROUNDELEMENT} \quad class_{fe}(b) = static$$
$$equal_{Background}(b_1, b_2) \equiv equal_{Universe}(b_1, b_2)$$
$$\forall b \in \text{BACKGROUNDELEMENT} \quad bkg(b) = \text{"Background"}$$

Rule Elements

ASM rules defined in a CoreASM specification (more precisely, defined in the current state of the simulated machine) are modeled by elements of the domain RULEELEMENT. States of CoreASM hold a mapping of rule names to rule elements defined in those states:

$$stateRule : \text{STATE} \times \text{NAME} \mapsto \text{RULE}$$
$$rules : \text{STATE} \mapsto \text{SET}(\text{RULE})$$
$$rules(s) \equiv \{r \mid r \in \text{RULE} \wedge (\exists n \in \text{NAME}, stateRule(s, n) = r)\}$$

Every rule element has a name[6], a body (which is a node of the parse tree) and a sequence of parameter names, all defined by the following functions:

$$name_{re} : \textsc{Rule} \mapsto \textsc{Name}$$
$$body : \textsc{Rule} \mapsto \textsc{Node}$$
$$param : \textsc{Rule} \mapsto \textsc{List}(\textsc{Name})$$

The equality of two rules is defined as the equality of their names, program bodies, and list of parameters.

$$equal_{Rule}(r_1, r_2) \equiv$$
$$name_{re}(r_1) = name_{re}(r_2) \wedge body(r_1) = body(r_2) \wedge param(r_1) = param(r_2)$$

For all $r \in \textsc{Rule}$, we have $bkg(r) =$ "Rule".

Enumerable Elements

In CoreASM, an element is called *enumerable* if it can be viewed as a collection (i.e., multiset) of other elements. The idea of enumerable elements provides a unique and yet simple interface to sets, multisets, lists, and other data structures. We define the following functions as the interface of enumerable elements:

- $enumerable : \textsc{Element} \mapsto \textsc{Boolean}$
  holds *true* if the element is enumerable. By default, $enumerable(e) = false$ for every element e unless otherwise specified.

- $enumerate : \textsc{Element} \mapsto \textsc{Multiset}(\textsc{Element})$
  provides a collection of elements representing the internal structure of the enumerable element.
  $$enumerate(e) \equiv enumerate_{bkg(e)}(e)$$

- $size : \textsc{Element} \mapsto \textsc{Number}$
  returns the size of this enumerable. For every enumerable element e, we have $size(e) = |enumerate(e)|$.

[6]The names of rule elements, universe elements, and function elements should all be unique in any given CoreASM state.

- *contains* : ELEMENT × ELEMENT ↦ BOOLEAN
  $$contains(e_1, e_2) \equiv \begin{cases} true, & \text{if } enumerable(e_1) \wedge e_2 \in enumerate(e_1) \\ false, & \text{otherwise.} \end{cases}$$

Among the elements we have defined so far, universe elements are enumerable (and so are the background elements). We have:

$$\forall u \in \text{UNIVERSEELEMENT} \quad enumerable(u) \wedge enumerate(u) = \{x | member_{ue}(u, x)\}$$

5.2 The Interpreter

The Interpreter evaluates an annotated parse tree and depending on the type of the root node, assigns a value, a location, or a multiset of update instructions to the root of the tree. The Interpreter interacts with the Abstract Storage in order to obtain values from the current state.

In this section we recall the ExecuteTree rule we presented in Section 4.2 and provide further details on the process of evaluating parse tree nodes. More specifically, this section refines the macro rule KernelInterpreter used by ExecuteTree.

5.2.1 Notation

We specify the Interpreter as a collection of rules (some embedded in the kernel, others contributed by plugins) which traverse a parse tree while evaluating values, locations and updates.[7] In order to introduce these rules, we state the following assumptions:

1. Nodes of the parse tree belong to the NODE universe and the following functions are defined on nodes:
   - *first* : NODE ↦ NODE, *next* : NODE ↦ NODE, *parent* : NODE ↦ NODE are static functions that implement tree navigation; by using these functions, the Interpreter can access all the children nodes of a given node, or access its parent (see Figure 4.3 for reference).
   - *class* : NODE ↦ CLASS returns the syntactical class of a node (i.e., the name of the corresponding grammar non-terminal class); for example RuleDeclaration .

[7] This section is a revised and extended version of what we have previously published in [48, Section 3].

- *grammarRule* : NODE $\mapsto$ GRAMMARRULE returns the grammar rule that produced that node.
- *token* : NODE $\mapsto$ TOKEN returns the syntactical token represented by the node (i.e., either a keyword, an identifier, or a literal value).
- *pattern* : NODE $\mapsto$ PATTERN returns the symbolic name for the specific grammar pattern corresponding to the node; for example, IfThen symbolically represents the pattern **if** ... **then**.
- $[\![\cdot]\!]$: NODE $\mapsto$ LOCATION $\times$ MULTISET(UPDATE) $\times$ ELEMENT holds the result of the interpretation of a node, given by a triple formed by a location (that is, the l-value of an expression, when it is defined), a multiset of update instructions, and a value (that is, the r-value of an expression)[8]. We access elements and establish properties of such triples through the following derived functions:
  - *loc* : NODE $\mapsto$ LOCATION returns the location (l-value) associated to the given node, i.e. $loc(n) \equiv [\![n]\!] \downarrow 1$.
  - *updates* : NODE $\mapsto$ MULTISET(UPDATE) returns the updates associated to the given node, i.e. $updates(n) \equiv [\![n]\!] \downarrow 2$.
  - *value* : NODE $\mapsto$ ELEMENT returns the value (r-value) associated to the given node, i.e. $value(n) \equiv [\![n]\!] \downarrow 3$.
  - *evaluated* : NODE $\mapsto$ BOOLEAN indicates if a node has been evaluated. We have,
$$evaluated(n) \equiv [\![n]\!] \neq undef$$
- *plugin* : NODE $\mapsto$ PLUGIN is the plugin associated to a node, that is, the plugin responsible for parsing and evaluation of the node.

2. A special variable *pos* holds at all times the current position in the tree (i.e., the current node being evaluated).

3. We use a form of pattern matching which allows us to concisely denote complex conditions on the nodes. In particular:
   - we denote with $\boxed{?}$ a generic node;

[8] The structure of the triple is intended to be mnemonic, with the l-value in the leftmost and the r-value in the rightmost position in the triple.

- we denote with $\boxed{\ }$ a generic unevaluated node; as an aid to the reader, we will also use the semantically equivalent $\boxed{e}$, $\boxed{r}$, and $\boxed{l}$ to denote unevaluated nodes whose evaluation is expected to result respectively, in a value (from an expression), a multiset of updates (from a rule), and a location;
- we denote with x an identifier node;
- we denote with v (value) an evaluated expression node (that is, a node whose *value* is not *undef*); we denote with u (update multiset) an evaluated statement node (a node whose *updates* is not *undef*); we denote with l (location) an evaluated expression for which a location has been computed (a node whose *loc* is not *undef*). We will at times add subscripts to these variables, or use different names for special cases that will be discussed as appropriate;
- we use prefixed Greek letters to denote positions in the parse tree (typically children of the current node, as denoted by *pos*) as in **if** $^{\alpha}e$ **then** $^{\beta}r$ where α and β denote, respectively, the condition node and the then-part node of an if statement;
- rules of the form

  $$(\!|\, pattern \,|\!) \rightarrow \quad actions$$

  are to be intended as

  if *conditions* **then** *actions*

  where the *conditions* are derived from the pattern according to the conventions above, as more formally specified in Table 5.1; in the action part of such a rule, an unquoted and unbound occurrence of l is to be interpreted as the *loc* of the corresponding node; an unquoted and unbound occurrence of v is to be interpreted as the *value* of the corresponding node; an unquoted and unbound occurrence of u as the *updates* of the corresponding node; and an unquoted and unbound occurrence of x as the *token* of the corresponding node.

Table 5.2 exemplifies how our compact notation can be translated into actual ASM rules.

4. The value of local variables (e.g., those defined in **import** and **let** rules) is maintained by a global dynamic function of the form env : TOKEN $\mapsto$ ELEMENT. We have

   $$env(x) \equiv top(envStack(x))$$

Abbreviation	Condition part	Action part
α, β etc.		*first*(pos), *next*(*first*(pos)), etc.
$^{\alpha}\boxed{?}$	$class(\alpha) \neq \mathsf{Id}$	
$^{\alpha}\boxed{\ }$	$class(\alpha) \neq \mathsf{Id} \wedge \neg evaluated(\alpha)$	
$^{\alpha}\boxed{e}$, $^{\alpha}\boxed{r}$, $^{\alpha}\boxed{l}$ *	$class(\alpha) \neq \mathsf{Id} \wedge \neg evaluated(\alpha)$	
$^{\alpha}x$	$class(\alpha) = \mathsf{Id}$	$token(\alpha)$
$^{\alpha}v$	$value(\alpha) \neq undef$	$value(\alpha)$
$^{\alpha}u$	$updates(\alpha) \neq undef$	$updates(\alpha)$
$^{\alpha}l$	$loc(\alpha) \neq undef$	$loc(\alpha)$

* These symbols are semantically equivalent to the $\boxed{\ }$ symbol; as a visual cue to the reader, the embedded letters express the intended result of evaluation.

Table 5.1: Abbreviations in Syntactic Pattern-matching Rules

where *envStack* is a function of the form *envStack* : TOKEN $\mapsto$ STACK(ELEMENT) which can be maintained by the following rules:

Interpreter

AddEnv$(x, v) \equiv$ Push$(envStack(x), v)$
RemoveEnv$(x) \equiv$ Pop$(envStack(x))$

Notice that, according to the rule ExecuteTree previously described in Section 4.2, interpreter rules in the kernel or from plugins are only executed when *evaluated*(pos) does not hold, i.e. when the current node has not been fully evaluated yet. Control moves from node to node either by explicitly assigning values to pos, or by setting $[\![pos]\!]$ to a value that is not *undef*; in which case, control is returned to the parent of pos by the ExecuteTree rule (unless an explicit assignment to pos is also made in the same step). Hence, the general strategy in our rules will be to evaluate all needed subtrees of a node, if any, by orderly assigning pos accordingly; when all needed subtrees are evaluated, we compute the resulting location, updates or value and assign it to $[\![pos]\!]$, thus implicitly returning control back to our parent. As exemplified in Table 5.2, our notation allows us to clearly visualize this process by the progressive substitution of evaluated u nodes for unevaluated $\boxed{r}$ nodes, and of v or l nodes for unevaluated $\boxed{e}$ nodes. Notice that identifiers do not have to be evaluated, hence we do not need a "boxed" version of x.

Compact notation	Actual rule
$(\!\vert$ **if** $^{\alpha}\boxed{e}$ **then** $^{\beta}\boxed{r}$ $\vert\!)\rightarrow$ $pos := \alpha$	**let** $\alpha = first(pos), \beta = next(first(pos))$ **in** **if** $class(pos) \neq \mathsf{Id}$ $\wedge\ pattern(pos) = \mathsf{IfThen}$ $\wedge\ class(\alpha) \neq \mathsf{Id}$ $\wedge\ \neg evaluated(\alpha)$ $\wedge\ class(\beta) \neq \mathsf{Id}$ $\wedge\ \neg evaluated(\beta)$ **then** $pos := first(pos)$
$(\!\vert$ **if** $^{\alpha}v$ **then** $^{\beta}\boxed{r}$ $\vert\!)\rightarrow$ **if** $v = \mathsf{true}_e$ **then** ...	**let** $\alpha = first(pos), \beta = next(first(pos))$ **in** **if** $class(pos) \neq \mathsf{Id}$ $\wedge\ pattern(pos) = \mathsf{IfThen}$ $\wedge\ value(\alpha) \neq undef$ $\wedge\ class(\beta) \neq \mathsf{Id}$ $\wedge\ \neg evaluated(\beta)$ **then** **if** $value(\alpha) = \mathsf{true}_e$ **then** ...
$(\!\vert$ **if** $^{\alpha}v$ **then** $^{\beta}u$ $\vert\!)\rightarrow$ ...	**let** $\alpha = first(pos), \beta = next(first(pos))$ **in** **if** $class(pos) \neq \mathsf{Id}$ $\wedge\ pattern(pos) = \mathsf{IfThen}$ $\wedge\ value(\alpha) \neq undef$ $\wedge\ updates(\beta) \neq undef$ **then** ...

Table 5.2: Examples of Pattern Matching Notation Translated into ASM Rules

5.2.2 Kernel Expression Interpreter

As previously described, the kernel interpreter rules implement the Boolean domain (but not the Boolean algebra), function evaluation and rule call (which share the same syntactic pattern), assignment, and import statement. We present in this section rules that result in values, namely for evaluating literals (true, false, undef) and nullary or n-ary functions.

Literals are simply lifted to their semantic counterparts:

Interpreter: Kernel Expressions

$(\!\|\,\textbf{true}\,\|\!)$	$\rightarrow$	$[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
$(\!\|\,\textbf{false}\,\|\!)$	$\rightarrow$	$[\![pos]\!] := (undef, undef, \mathsf{false}_e)$
$(\!\|\,\textbf{undef}\,\|\!)$	$\rightarrow$	$[\![pos]\!] := (undef, undef, \mathsf{undef}_e)$
$(\!\|\,\textbf{self}\,\|\!)$	$\rightarrow$	$[\![pos]\!] := (undef, executingAgent, \mathsf{undef}_e)$

Evaluation of identifiers as expressions depends on whether the identifier refers to a local variable or a function. To evaluate an identifier as an expression, the Interpreter first checks the set of in-scope local variables for a possible value for the identifier. If the identifier was not a local variable (i.e., it is not found in the local environment), the Interpreter checks if the identifier refers to a (nullary) function, in which case the Abstract Storage is queried for the value of that function in the current state. If instead the identifier is not defined, the macro HandleUndefinedIdentifier (described later) is called. The rule for n-ary functions is similar, except that the arguments of the function are evaluated first. The formal definition is as follows:

Interpreter: Kernel Expressions

$(\!|\,^{\alpha}x\,|\!) \rightarrow$

if $env(x) \neq undef$ **then**
  $[\![pos]\!] := (undef, undef, env(x))$
else
  if $isFunctionName(x)$ **then**
    let $l = (x, \langle\rangle)$ **in**
      $[\![pos]\!] := (l, undef, getValue(l))$
  if $undefinedToken(x)$ **then**
    HandleUndefinedIdentifier$(pos, x, \langle\rangle)$

$(\!|\, {}^{\alpha}x({}^{\lambda_1}\boxed{?}_1, \ldots, {}^{\lambda_n}\boxed{?}_n) \,|\!)$ → **if** *isFunctionName*(x) **then**
choose $i \in [1..n]$ **with** $\neg$*evaluated*(λ_i)
$pos := \lambda_i$
ifnone
let $l = (x, \langle value(\lambda_1), \ldots, value(\lambda_n)\rangle)$ **in**
$[\![pos]\!] := (l, undef, getValue(l))$
if *undefinedToken*(x) **then**
HandleUndefinedIdentifier($pos, x, \langle\lambda_1, \ldots, \lambda_n\rangle$)

where
undefinedToken(x) $\equiv \neg$(*isFunction*(x) $\vee$ *isRule*(x) $\vee$ *isUniverse*(x))

Notice how in the second pattern, the $\boxed{?}$ symbol is used to denote arguments, both unevaluated and evaluated. If x is bound to a function, the rule specifies that all arguments must be evaluated, without any specific order, to determine the location of the node. While there are still unevaluated arguments, the rule sets *pos* to the node representing an unevaluated argument; as soon as the evaluation of the argument is complete, control returns to the parent node (and thus, again to the same rule), until all arguments are evaluated. At this point (**ifnone** branch), the location and values of the function are computed and stored in $[\![pos]\!]$.

Finally, if the Interpreter encounters an identifier that is not bound to any element of the state, the HandleUndefinedIdentifier rule (see Appendix A.2) will consult all the plugins that are registered to handle undefined identifiers. More specifically, such plugins are asked to evaluate the node with the undefined identifier.[9] If none of the plugins could evaluate the node, KernelHandleUndefIdentifier will be called to create a new function element with a default value of undef_e for the given arguments. This default behavior of the kernel is a "liberal" approach toward type-checking; it allows identifiers to be used without declaration, which is suited for early analysis and specification.

Interpreter: Undefined Identifier

KernelHandleUndefIdentifier($pos, x, args$) ≡
let $f = new$(FUNCTIONELEMENT) **do**
stateFunction($state, x$) := f
$[\![pos]\!] := ((x, args), undef, \mathsf{undef}_e)$

---

[9] It is considered an error if more than one plugin evaluate the undefined identifier with different results.

5.2.3 Kernel Rule Interpreter

Rule plugins provide the execution semantics of rules. Execution of rules results in a multiset of update instructions that is the underlying value for the rule node of the parse tree. As discussed in Section 4.2, accumulated update instructions are used by the Abstract Storage to compute the updates set that will ultimately be applied to the current state to generate the next state.

We start with the **skip** rule or the no-operation rule. The semantics of the **skip** rule is simply to produce an empty multiset of updates:

Interpreter: Kernel Rules

$(\!|\,\textbf{skip}\,|\!) \;\rightarrow\; [\![pos]\!] := (undef, \{\!|\}, undef)$

Rule Calls

To evaluate an identifier as a rule, the Interpreter first checks if a rule element is bound to the identifier. If so, the RuleCall macro is called to execute the rule. Notice that in this case, arguments are *not* evaluated prior to calling the rule: in fact, the semantics of rule calls in [25] prescribes that the formal parameter in the body of the rule must be substituted with the entire term that is used as the actual argument, not its value.

Interpreter: Kernel Rules

$(\!|\,{}^{\alpha}x\,|\!) \;\rightarrow\;$ **if** $isRuleName(x)$ **then** RuleCall$(ruleValue(x), \langle\rangle)$

$(\!|\,{}^{\alpha}x(^{\lambda_1}\boxed{?}_1, \ldots, ^{\lambda_n}\boxed{?}_n)\,|\!) \;\rightarrow\;$ **if** $isRuleName(x)$ **then** RuleCall$(ruleValue(x), \langle\lambda_1, \ldots, \lambda_n\rangle)$

Traditionally, rule calls in ASMs have been used in two form: as macros, or as submachines. The difference between the two forms is that calling a macro simply means executing its body (possibly with parameter substitution) and collecting the resulting updates, whereas running a submachine results in an entire encapsulated computation of the rule, that is iterated until completion, as defined in [25, Section 4.1.2]. Here, we model macro calls, while the effect of submachine calls can simply be achieved by using the **iterate** construct; see Section 6.1.8 for the specification of the **iterate** construct.

As we have already noted, ASMs differ from many other languages in that *call-by-substitution* is used for parameters instead of the more usual *call-by-value*. In other words, actual parameters are evaluated at the point of use (in the callee) rather than at the point of call (in the caller). Due to the presence of **seq**-rules, the difference can be observable, as parameters can be evaluated in different states. Hence, we have to substitute the whole parse tree denoting an actual parameter (i.e., an expression) for each occurrence of the corresponding formal parameter in the body of the callee. Also, we substitute parameters in a copy of the callee body, to avoid modifying the original definition.

There are several static semantic constraints on valid rule declarations; for example, it is assumed that the formal parameters of a rule are all pairwise distinct, and that the formal parameters are the only freely occurring variables in the body of the rule (see [25], Definition 2.4.18). For simplicity, we do not explicitly check for such conditions in our specification.

The RuleCall routine, defined below, describes how rule calls (possibly with parameters) are handled.

Interpreter: Kernel Rules

RuleCall$(r, args) \equiv$
  if $workCopy(pos) = undef$ **then**
    let $b' =$ CopyTreeSub$(body(r), param(r), args)$ **in**
      $workCopy(pos) := b'$
      $parent(b') := pos$
      $pos := b'$
  else
    $[\![pos]\!] := (undef, updates(workCopy(pos)), value(workCopy(pos)))$
    $workCopy(pos) := undef$

The rule CopyTreeSub returns a copy of the given parse tree, where every instance of an identifier node in a given sequence (formal parameters) is substituted by a copy of the corresponding parse tree in another sequence (actual parameters). We assume that the elements in the formal parameters list are all distinct (i.e., it is not possible to specify the same name for two different parameters). Also, formal parameters substitution is applied only to occurrences of formal parameters in the original tree passed as argument, and *not* also on the actual parameters themselves. See Appendix A.2 for the definition of CopyTreeSub.

Assignment and Import

The kernel of the CoreASM engine also includes assignment and **import** rules. Assignment is performed as follows:

Interpreter: Kernel Rules

$(\!| {}^{\alpha}\boxed{?} := {}^{\beta}\boxed{?} |\!)$ $\rightarrow$

choose $\tau \in \{\alpha, \beta\}$ **with** $\neg evaluated(\tau)$
  $pos := \tau$
ifnone
  if $loc(\alpha) \neq undef$ **then**
    if $isModifiable(stateFunction(state, name_{lc}(loc)))$ **then**
      $[\![pos]\!] := (undef, \{\!|\langle loc(\alpha), value(\beta)\rangle|\!\}, undef)$
    else
      Error('Cannot update a non-modifiable function')
  else
    Error('Cannot update a non-location.')

It is worthwhile to remark that the rule above does not syntactically constrain assignment to be performed exclusively to variables or functions: rather, any plugin can contribute new forms of expressions which, as long as they result in a modifiable location (e.g., not a monitored function), are deemed syntactically acceptable in the lhs of an assignment.

The **import** rule is defined as follows:

Interpreter: Kernel Rules

$(\!|$ **import** ${}^{\alpha}x$ **do** ${}^{\beta}\boxed{r}$ $|\!)$ $\rightarrow$

let $e = new(\text{ELEMENT})$ **in**
  AddEnv(x, e)
$pos := \beta$

$(\!|$ **import** ${}^{\alpha}x$ **do** ${}^{\beta}u$ $|\!)$ $\rightarrow$

RemoveEnv(x)
$[\![pos]\!] := (undef, u, undef)$

To perform an **import**, a new element is created and it is assigned to the value of the given identifier (x) in the local environment. The rule part $\boxed{r}$ is then evaluated in this new environment by assigning *pos* to the corresponding node. The identifier is then removed from the local environment when the evaluation of the rule part is complete.

5.2.4 Operators

Although plugins can extend the CoreASM language by introducing (almost) arbitrary expression forms, operators are treated specially in the CoreASM engine. To avoid lengthy expressions with unnecessary parenthesis, the engine provides plugins with a mechanism to declare a precedence level for the operators they contribute.

Precedence level of an operator is defined by a numeric value $p \in [0 \ldots 1000]$, where 1000 is the highest priority. This value should be attached to all operator patterns. The following example introduces a new operator Ω with precedence level 300:

$$(\!| \boxed{e} \ \ \Omega \ \boxed{e} \ |\!)_{[300]} \rightarrow \ \ldots$$

The only operator provided by the kernel is the equality operator ("="). Two values are considered to be equal if they are equal according to at least one of their corresponding backgrounds. In the following rule, the equality functions provided by the backgrounds of the operands are queried to determine the equality:

Interpreter: Kernel Operators

$(\!| {}^{\alpha}\boxed{?} = {}^{\beta}\boxed{?} |\!)_{[600]} \rightarrow$

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  let $e_1 = value(\alpha),\ e_2 = value(\beta)$ **in**
    let $b_1 = bkg(e_1),\ b_2 = bkg(e_2)$ **in**
      if $equal_{b_1}(e_1, e_2) \vee equal_{b_2}(e_2, e_1)$ **then**
        $[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
      else
        $[\![pos]\!] := (undef, undef, \mathsf{false}_e)$

5.3 Rules and Updates

According to the original definition of ASMs, evaluation of each ASM rule results in a potentially empty set of updates of the form (l, v) where l is a location and v is a value to be assigned to that location if the update is successfully applied to the state. At the end of each computation step, the update set produced by evaluating the program of the machine (or programs of the agents in a multi-agent ASM), if consistent, will be applied to the state

to form the new state.[10]

In CoreASM, we originally followed exactly the same idea: rules would produce update sets of the form $\langle l, v\rangle$. However, this approach would seriously limit and complicate incremental or partial modification of elements with internal structure that are composed of other elements, such as sets, maps, and trees. For example, parallel addition of elements 5 and 7 to the set $\{1, 2\}$ residing at the location $f(a)$ would lead to two inconsistent updates of $\langle(\text{“f”}, \langle a\rangle), \{1, 2, 5\}\rangle$ and $\langle(\text{“f”}, \langle a\rangle), \{1, 2, 7\}\rangle$.

Inspired by the idea of *partial updates* introduced in [76, 77], we extend CoreASM updates from a pair of location-value to a triplet of the form

$$\langle \textsc{Location}, \textsc{Element}, \textsc{Action}\rangle,$$

called *update instruction*, that is general enough to represent regular and partial updates.[11] Update instructions consist of a location, a value, and an *action* that defines the type of modification that has to be done on the location. The most basic action, which is defined in the CoreASM kernel, is the *updateAction* $\in$ Action. An update instruction of the form $\langle l, v, \textit{updateAction}\rangle$ is semantically equivalent to an original ASM update of (l, v). However, background plugins may introduce their own special actions; for example, a plugin providing the background of sets may introduce two new actions *setAddAction* and *setRemoveAction* respectively representing the actions of adding and removing elements from a set. As a result, in our example of adding 5 and 7 to the set $\{1, 2\}$ above, the parallel execution of the rules will lead to the following update instructions: $\langle(\text{“f”}, \langle a\rangle), 5, \textit{setAddAction}\rangle$ and $\langle(\text{“f”}, \langle a\rangle), 7, \textit{setAddAction}\rangle$ which will have to be later *aggregated* by the plugin into one single regular update on the given location.

5.3.1 Update Instruction Notation

We define the following functions on update instructions:

- *uiLoc* : Update $\mapsto$ Location
  returns the location associated with the given update instruction.

[10] The ideas presented in this section has been previously discussed in more detail in M. Memon's M.Sc. thesis [99].

[11] In practice, we define update instructions as quadruples of the form $\langle$Location, Element, Action, Set(Element)$\rangle$ where the 4^{th} element is the set of agents that produced the update instruction (an update may be the result of aggregating two or more updates); however, in this work we often leave out the reference to the 4^{th} element and view update instructions as triples.

- *uiVal* : UPDATE $\mapsto$ ELEMENT
  returns the value associated with the given update instruction.

- *uiAction* : UPDATE $\mapsto$ ACTION
  returns the action associated with the given update instruction.

- *uiAgents* : UPDATE $\mapsto$ SET(ELEMENT)
  returns the set of agents that produced the given update instruction.

- *aggStatus* : UPDATE $\times$ PLUGIN $\mapsto$ {*successful*, *failed*}
  indicates the aggregation status of an update instruction, set by a given aggregator plugin. If an update instruction *ui* has not been processed by a plugin, *aggStatus*(*ui*) is *undef*.

5.3.2 Aggregation of Updates

According to the original ASM definition, after every computation step, location contents are changed by and only by updates. In order to be faithful to that definition, with the introduction of partial updates, we introduce an *aggregation phase* in every computation step that takes place before the application of updates to the state. *Aggregation* is the process of combining all update instructions affecting a single location, into one single update which is called the *resultant update*. The aggregation phase of a CoreASM step performs aggregation on all locations affected by the step and results in a set of regular updates.[12]

Since the CoreASM kernel does not introduce any special update actions other than the one for regular updates, it only defines the framework in which background plugins can provide their background-specific partial updates and their corresponding aggregation algorithms. We say that a plugin is *responsible* for an action, if it is registered to aggregate update instructions of that action. A plugin is said to be *responsible* for aggregation of a given update instruction if the update instruction contains an action for which the plugin is responsible. Finally a plugin is considered to be *responsible* for aggregation of a given regular update if there is an update instruction that operates on the the same location. A plugin that is registered for aggregation of one or more update action is called an *aggregator* plugin.

[12] This is also in line with the *integration* phase introduced in [76].

Recalling the definition of AggregateUpdates on page 59, Abstract Storage calls the following rule in its *Aggregation* control state before firing the updates to the state (see also Figure 4.8):

Abstract Storage

AggregateUpdates ≡
  updateSet ← Aggregate(*updateInstructions*)

The Aggregate method runs the aggregation method of all the aggregator plugins on the update instructions, gathers the resulting updates and returns the compiled set. When called for aggregation, an aggregator plugin aggregates all update instructions for which it is responsible and flags them as either successful or failed. It is important to note that the order in which plugins are called to perform aggregation should not affect the resultant updates produced. Also note that the failure in aggregation of a single plugin should not fail the aggregation attempt of other plugins.

Abstract Storage

Aggregate(*updates*) ≡
  let $ap = \{a \mid a \in \text{PLUGIN} \wedge aggregator(a)\}$ **in**
    seq
      forall $p \in ap$ **do**
        let $R = aggregatorRule(p)$ **in**
          $resultantUpdates(p, updates) \leftarrow R(updates)$
    next
      result := $\bigcup_{p \in ap} resultantUpdates(p, updates)$

The *resultantUpdates* function is used to collect resultant updates from plugins for a given multiset, and the $aggregatorRule(p)$ function returns the aggregation rule provided by plugin p. Note that a plugin aggregator rule is expected to accept a multiset of update instructions as an argument, and its invocation should cause the return of its resultant updates with the return-result rule syntax as described in [25, Def. 4.1.7].

Plugin Aggregation Consistency

Aggregation algorithms provided by plugins also implicitly define the acceptable semantics of the combination of updates they process. During an aggregation process, a plugin may encounter a situation where the updates and instructions for a given location cannot be

aggregated into a regular update. Such a situation may occur, for example, if there are updates or instructions that are semantically inconsistent, such as addition and removal of the same element from a set.

When the aggregation of all updates and instructions affecting a given location are deemed inconsistent, the plugin flags all updates to the location as *failed*.

Abstract Storage

HandleInconsistentAggregation$(loc, updateMset, plugin) \equiv$
  forall $ui \in updateMset$ **with** $uiLoc(ui) = loc$ **do**
    $aggStatus(ui, plugin) := failed$

Although aggregation for a single location may have failed, the aggregation of the rest of the update instructions a plugin is responsible for would continue.

Basic Update Aggregator

Once aggregation of all aggregator plugins have completed successfully, the resultant update set may still have updates with a regular update action that do not need aggregation but are not flagged as processed. The *Basic Update Aggregator* provided by the Kernel plugin (see Section 4.2.1) solves this problem by returning a set of all regular updates for locations which do not require any aggregation and flagging all those updates as *successful*. The basic update aggregator is called by AggregateUpdates along side all aggregator plugins.

Abstract Storage

BasicUpdateAggregator$(updateMset) \equiv$
  seq
    result $:= \{\}$
  next
    forall $ui \in updateMset$ **with** $uiAction(ui) = updateAction$ **do**
      if $\nexists\, ui2 \in updateMset,\ uiLoc(ui) = uiLoc(ui2) \wedge uiAction(u2) \neq updateAction$ **then**
        add ui **to result**
        $aggStatus(ui, kernelPlugin) := successful$

5.3.3 Composition of Updates

Aggregation as we have described it so far gives semantically acceptable results with basic ASMs. However, for Turbo ASMs, which allow for sequential composition and iteration

of ASMs within one single step of the machine, aggregation alone is insufficient. While the sequential composition of ASMs imposes an order between the sets of updates (on a location), it is not always desirable for a Turbo ASM rule to return aggregated resultant updates. On the other hand, update instructions produced by a Turbo ASM rule has to be composed in a form that preserves the sequential semantics of the updates. As an example, consider the following sequential composition, where $s = \{1, 2\}$:

seq
  add 5 **to** s
  add 7 **to** s
next
  remove 5 **from** s
  add 6 **to** s

The semantics of this rule is to add 6 and 7 to s. Since this rule may be executed in parallel with other rules that may also modify the set s, it is desirable that the evaluation of this rule does not result in aggregated updates (i.e., a regular update assigning $\{1, 2, 6, 7\}$ to s). On the other hand, there is an explicit order between the update instructions produced by the two parts of this sequence which has to be reflected in the resulting update multiset. As a result, a special *composition* process has to be introduced on update instructions that composes two multisets of update instructions into one multiset with respect to the order of updates. In the above example, removing 5 from s neutralizes the addition of 5 in the first step and so neither of the two modifications will appear in the result of the composition, which will be $\{\!|\langle(\text{“s”}, \langle\rangle), 6, \mathit{setAddAction}\rangle, \langle(\text{“s”}, \langle\rangle), 7, \mathit{setAddAction}\rangle|\!\}$.

Since the CoreASM kernel does not define any special update action, its composition (captured by the Compose rule defined below) basically relies on the composition behaviors provided by background plugins. As a result, every aggregator plugin is required to also provide a composition algorithm which, when given two update multisets, produces composed update instructions for all locations for which the plugin is responsible.

It is important to note that the Compose rule expects the first update multiset to be consistent with respect to typical ASM consistency and aggregation consistency. The result of sequential composition of the two update multisets would then be the union of all composed update instructions produced by individual plugins.

Abstract Storage

Compose$(uMset_1, uMset_2) \equiv$
  seq
    let $ap = \{a \mid a \in \text{PLUGIN} \wedge aggregator(a)\}$ **in**
      forall $p \in ap$ **do**
        let $R = composerRule(p)$ **in**
          $composedUpdates(p, uMset_1, uMset_2) \leftarrow R(uMset_1, uMset_2)$
  next
    result $:= \bigcup_{p \in ap} composedUpdates(p, uMset_1, uMset_2)$

In the above rule, the *composedUpdates* function is used to collect the updates resulting from plugins performing sequential composition of two update multisets. The *composerRule* function is expected to return the composition behavior of the given plugin, implementing the composition of updates on locations for which it is responsible. Note that the composition rule for each plugin is expected to accept two multisets as arguments, and its invocation should cause the return of the sequentially composed update multiset with the return-result rule syntax as described in [25, Def. 4.1.7].

A plugin which provides aggregation, must also provide facilities for sequential composition of actions for which it is responsible. A plugin is deemed responsible for the composition of updates at a given location, if and only if:

- The plugin is responsible for aggregation of the location with respect to the second update multiset.
- The plugin is responsible for aggregation of a location with respect to the first update multiset, if and only if that location is not affected by the second update multiset.

Basic Update Composer

To complement the basic update aggregator we introduced earlier, the Kernel plugin also provides a default update composition behavior. The *Basic Update Composer* is responsible for performing sequential composition of locations affected solely by basic updates. Sequential composition of updates in basic ASMs (without partial updates) is formally defined in [25, Def. 4.1.1] as

$$U \oplus H = \{u \in U \mid location(u) \notin locations(H)\} \cup H$$

In CoreASM, with the existence of partial updates, sequential composition of basic updates is similarly defined as:

$$compose(U, H) \equiv \{u \in U \mid location(u) \notin locations(H) \wedge isBasicUpdate(u)\} \cup \{u \in H \mid isBasicUpdate(u)\}$$

The basic update composer is then defined as follows:

Abstract Storage

BasicUpdateComposer$(uMset_1, uMset_2) \equiv$

result $:= \{ui_1 \mid ui_1 \in uMset_1 \wedge isBasicUpdate(uMset_1, ui1) \wedge \neg locUpdated(uMset_2, uiLoc(ui_1))\}$
$\cup \{ui_2 \mid ui_2 \in uMset_2 \ \wedge isBasicUpdate(uMset_2, ui_2)\}$

where

$isBasicUpdate(uMset, ui) \ \equiv \ \forall \langle l, v, a\rangle \in uMset, \ l = uiLoc(ui) \Rightarrow a = updateAction$

$locUpdated(uMset, l) \ \equiv \ \exists \ ui \in uMset, \ uiLoc(ui) = l$

We refer to Mashaal Memon's M.Sc. thesis [99] for further details on aggregation and composition of updates.

5.4 The Parser

CoreASM offers the possibility of extending and modifying the syntax and semantics of its language, keeping only the bare essential parts of the ASM language as static. In order to achieve this goal, CoreASM plugins should be able to extend the grammar of the core language by providing new grammar rules together with their semantics. As a result, the kernel of the engine does not have a comprehensive parser. Plugins used in a given specification can provide portions of the grammar (sets of grammar rules) of the language based on which the specification has to be parsed. Upon loading a specification, the engine will combine all the provided grammar rules into a single grammar. Based on this grammar, a parser is generated which will be used to generate the parse tree of the specification. Hence, the CoreASM parser is in fact a *parser generator* which, when given a grammar, produces a parser that can be used to parse a given specification. As a result, the grammar used for two different specifications may be different, depending on the plugins required by the specifications. One of the challenges in the implementation of CoreASM had been to equip the engine with a fast parser generator capable of generating parsers with look-ahead of

more than one to allow co-existence of more than one grammar rule starting with the same pattern.

We do not intend to specify the details of the CoreASM parser; we only require that the parser provides the following function and rule as part of its interface:

- A function of the form *requestedPlugins* : SPECIFICATION $\mapsto$ SET(PLUGIN) that for every specification returns the list of plugins used by that specification. In practice, this would be achieved by looking for the **use** clauses in the specification.
- An ASM rule of the form Parse($spec, \mathcal{G}$) that parses the given specification *spec* with respect to the given grammar $\mathcal{G}$, produces a parse tree of nodes (values of the domain NODE, see Section 5.2.1) representing the specification, and returns the root node of the parse tree.

5.5 The Plugin Framework

The CoreASM plugin architecture supports two extension mechanisms: plugins can either extend the functionality of specific components of the engine, by contributing additional data or behavior to those components (i.e., adding new grammar rules to the Parser, new semantic rules to the Interpreter, new backgrounds, universes, and functions to the Abstract Storage, and new policies to the Scheduler) or they can extend the control state ASM of the engine, by interposing their own code in between state transitions.

Practically speaking, a CoreASM plugin can be implemented as a Java class that implements one or more of the interfaces defined by the CoreASM extensibility framework (see Table 5.3 and also Section 7.2.1). In this section we look at various plugin interfaces and explore the mechanisms through which they extend the CoreASM engine.

5.5.1 Parser Extensions

Plugins can implement the *Parser Plugin* interface and/or the *Operator Provider* interface to extend the Parser by respectively contributing additional grammar rules and new operator descriptions. We assume that for any parser plugin *pp*, *pluginGrammar*(*pp*) holds the set of all the grammar rules contributed by *pp*, and for any operator provider *op*, *pluginOperators*(*op*) holds the descriptions (syntax and semantics) of new operators contributed by *op*.

Plugin Interface	Extends	Description
Parser Plugin	Parser	provides additional grammar rules to the parser
Interpreter Plugin	Interpreter	provides new semantics to the Interpreter
Operator Provider	Parser, Interpreter	provides grammar rules for new operators along with their precedence levels and semantics
Vocabulary Extender	Abstract Storage	extends the state with additional functions, universes, and backgrounds
Aggregator	Abstract Storage	aggregates partial updates into basic updates
Scheduler Plugin	Scheduler	provides new scheduling policies for multi-agent ASMs
Extension Point Plugin	all components	extends the control state model of the engine

Table 5.3: CoreASM Plugin Interfaces

Before parsing a specification, the engine gathers all the grammar rules and operator descriptions provided by all parser plugins and operator providers. The Parser then combines these grammar rules and operator descriptions with the kernel grammar and builds a new 'parser' to scan the specification. While building the abstract syntax tree, this parser labels the nodes that are created by plugin-provided grammar rules with the plugin's identifier; these labels can later be used by the Interpreter to evaluate the nodes.

Parser plugins and operator providers are probed by the LoadSpecPlugins rule before the engine starts parsing the specification (see Figure 4.5). This rule iterates over all the plugins required by the loaded specification and after ensuring dependency requirements, loads the plugins by calling the LoadPlugin rule presented below. The latter initializes the plugin, then loads all the provided grammar rules and operator descriptions to be processed by the parser in the next step of the process.

Control API

```
LoadPlugin(p) ≡
  if p ∉ loadedPlugins then
    seq
      InitializePlugin(p)
    next
      add p to loadedPlugins
      if isParserPlugin(p) then
        add pluginGrammar(p) to grammarRules
      if isOperatorProvider(p) then
        add pluginOperators(p) to operatorRules

InitializePlugin(p) ≡
  let R = pluginInitRule(p) in
    R
```

5.5.2 Interpreter Extensions

Plugins can extend the Interpreter component of the engine by implementing either the *Interpreter Plugin* interface or the *Operator Provider* interface (or both). These plugins provide the semantics for rules and operations contributed as per Section 5.5.1. Traversing the abstract syntax tree, the ExecuteTree rule of the Interpreter (see Figure 4.9) uses these semantic rules to evaluate nodes that correspond to the extended grammar rules.

The semantics contributed by a plugin p which implements the Interpreter Plugin interface can be obtained through *pluginRule*(p). As already mentioned earlier, nodes of the parse tree corresponding to grammar rules provided by a plugin are annotated with the plugin identifier. If a node is found to refer to a plugin, the Interpreter obtains the semantic rules provided by that plugin and executes it; otherwise, the default kernel Interpreter rules are used (see ExecuteTree on page 58).

A similar approach is also used by the KernelInterpreter rule to obtain semantics of extended operators from operator providers. A detailed discussion on how the engine deals with operators and their extensions is provided in [99].

5.5.3 Abstract Storage Extensions

Vocabulary Extender plugins extend the vocabulary of the CoreASM state by contributing new backgrounds, universes, and functions to the Abstract Storage. Such plugins in fact extend the initial state and the signature of the simulated machine. The following functions, defined on vocabulary extender plugins, respectively hold the backgrounds, universes, functions, and rule elements such plugins provide:

$$pluginBackgrounds : \textsc{Plugin} \mapsto (\textsc{Name} \mapsto \textsc{BackgroundElement})$$
$$pluginUniverses : \textsc{Plugin} \mapsto (\textsc{Name} \mapsto \textsc{UniverseElement})$$
$$pluginFunctions : \textsc{Plugin} \mapsto (\textsc{Name} \mapsto \textsc{FunctionElement})$$
$$pluginRules : \textsc{Plugin} \mapsto (\textsc{Name} \mapsto \textsc{Rule})$$

In the Abstract Storage, *stateUniverse* and *stateFunction* bind the names of functions and universes in the CoreASM state to the mathematical objects that represent them (see Section 5.1). Backgrounds are considered as special universes and hence are handled by *stateUniverse*. The value of these functions is initialized by the InitAbstractStorage rule (see Figure 4.5). While creating the default universe and functions, the engine calls LoadVocabularyPlugins to iterate over all vocabulary extender plugins and to extend the CoreASM state with the vocabulary they provide.

Abstract Storage

LoadVocabularyPlugins(*state*) ≡
  forall $p \in specPlugins$ **do**
    if $isVocabularyExtender(p)$ **then**
      forall $(bkgName, bkg) \in pluginBackgrounds(p)$ **do**
        $stateUniverse(state, bkgName) := bkg$
      forall $(uName, universe) \in pluginUniverses(p)$ **do**
        $stateUniverse(state, uName) := universe$
      forall $(fName, f) \in pluginFunctions(p)$ **do**
        $stateFunction(state, fName) := f$
      forall $(rName, rBody) \in pluginRules(p)$ **do**
        $stateRule(state, rName) := rBody$

Plugins can also implement the *Aggregator* interface and provide aggregation and composition rules to be applied on update instructions before they are submitted to the state. Aggregator plugins are called to aggregate update instructions by the AggregateUpdate rule in the *Aggregation* state of the engine; see Figure 4.8 and Section 5.3.2 for more details. For

any aggregator plugin *ap*, *aggregatorRule*(*ap*) and *composerRule*(*ap*) respectively hold the aggregation and composition behaviors provided by *ap*.

5.5.4 Scheduler Extensions

Policy plugins, also called *Scheduler plugins*, extend the scheduler of the engine by providing new scheduling policies that affect the selection of agents in multi-agent ASMs. They provide an extension to the scheduler that is used to determine at each step the next set of agents to execute. We assume that for any scheduling plugin *sp*, *pluginSchedulingPolicy*(*sp*) holds the scheduling policy provided by *sp*. For any scheduling policy, the following functions should be defined:

- *newSchedulingGroup* : SchedulingPolicy $\mapsto$ SchedulingGroup
  returns a new scheduling group for the given policy. A scheduling group binds a group of schedules together. The exact semantics of such a group would be defined by the scheduling policy. For example, in a one-by-one scheduling policy that tries to offer a fair schedule, all the schedules created within a group share the same 'memory', i.e. they avoid scheduling already scheduled elements before scheduling the 'remaining' elements.

- *newScheduleRule* : SchedulingPolicy $\mapsto$ Rule
  returns an ASM rule modeling a function of the form

  $$f : \text{SchedulingGroup} \times \text{Set} \mapsto \text{List}(\text{Set})$$

  that given a scheduling group and an initial set of elements (agents), provides a new schedule based on the given policy. The schedule is in form of a list of subsets of the initial set of elements. For example, a schedule on the set $\{a, b, c\}$ can be $\langle\{a, b, c\}, \{a, b\}, \{b, c\}\rangle$ or $\langle\{c\}\rangle$.

See Section 6.4.2 for an example of a policy plugin.

5.5.5 Extension Point Plugins

In addition to modular extensions of specific components, plugins can also extend the control state of the engine by registering themselves for *Extension Points*. Each control state transition in the execution engine is associated to an extension point. At each extension

point, if there is any plugin registered for that point, the code contributed by the plugin for that transition is executed before the engine proceeds to the next control state. Such a mechanism enables arbitrary extensions to the engine's lifecycle, which facilitates implementing various practically relevant features such as adding debugging support, adding a C-like preprocessor, or performing statistical analysis of the behavior of the simulated machine (e.g., coverage analysis or profiling). A plugin, for example, could monitor the updates that are generated by a step before they are actually applied to the current state of the simulated machine, possibly checking conditions on these updates and thus implementing a kind of watches (i.e., displaying updates to certain locations) or watch-points (i.e., suspending execution of the engine when certain updates are generated), which are useful for debugging purposes. As an additional example, a plugin could provide syntax for declaring assertions and invariants. Assertions have to be checked when the corresponding node is evaluated, hence the plugin would also implement the Interpreter extension to give semantics to assertions. In contrast, invariants have to be checked at each step (not when a particular rule is executed), for example immediately before applying updates: thus, the plugin would hook on the FireUpdateSet extension point to check that the declared invariants really hold in each state.

As we mentioned earlier, we have used a variant of control state ASMs to present a high-level specification of the CoreASM engine. Recalling the definition of control state ASMs from Section 2.3, a control state ASM is an ASM whose rules are all of the form presented in Figure 2.1.

To model the CoreASM engine, we introduce a variation of control state ASMs, called an *Extensible Control State ASM*, which is a control state ASM with an additional (and potentially dynamic) set of *extension point plugins* contributing supplementary rules that are executed before the machine switches to a new state (i.e. before *ctl_state* gets a new value).

Extensible control state ASMs are pictured with almost the same control state diagrams as shown in Figure 2.1. The difference is that in EFSM diagrams, the transition with an extension point is marked with a small diamond;[13] see Figure 5.2(a) for an example. Rules of extensible control state ASMs are formulated in textual form by a set of *Extensible Finite State Machine* (EFSM) rules, where EFSM is defined as follows:

[13]In order not to confuse the reader, we have omitted the diamond from our diagrams. However, this should not be a concern since the extension points are always on the transitions leading to control states.

EFSM

```
EFSM(i, if cond then rule, j) ≡
  if ctl_state = i and cond then
    rule seq Proceed(i, j)

Proceed(i, j) ≡
  seq
    forall p ∈ extensionPointPlugins do
      marked(p) := isPluginRegisteredForTransition(p, i, j)
  seq
    iterate
      let eps = {p | p ∈ extensionPointPlugins with marked(p)} in
        choose p' ∈ eps with ∀p'' ∈ eps holds priority(p') ≥ priority(p'') do
          marked(p') := false
          let R = pluginExtensionRule(p') in
            R(i, j)
  next
    ctl_state := j
where
  priority(p) ≡ pluginCallPriority(p, i, j)
```

An EFSM rule, instead of updating the control state of the machine in parallel with the execution of the transition rule, first executes the transition rule, then iterates over all the extension point plugins (according to their priority) and one by one executes their extension rules before switching the control state of the machine to a new state.[14]

As an example, the extensible control state ASM of Figure 5.2(a) can be executed with a set of extension point plugins $\{p_1, p_2\}$ contributing rules $PRule_1$ and $PRule_2$ which extend the control state of the machine (during its execution) to the control state ASM of Figure 5.2(b).

The following functions are defined on extension point plugins:

- *isPluginRegisteredForTransition* : Plugin × EngineMode × EngineMode ↦ Boolean holds true if the given plugin is registered to extend the behavior of the transition between the two given engine modes.

[14] If two plugins have the same call priority, their rules will be executed in a non-deterministic order.

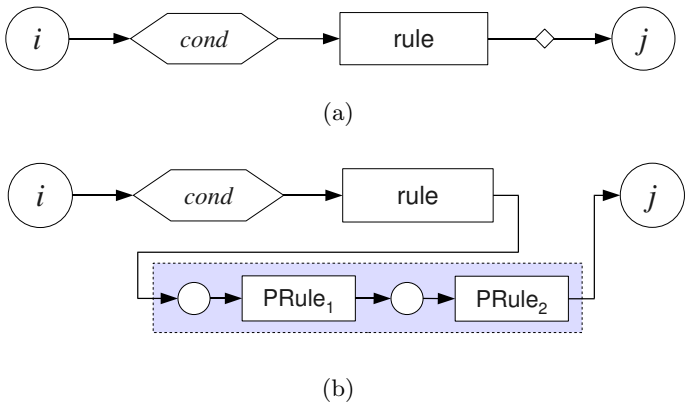

Figure 5.2: (a) An extensible control state ASM and (b) one of its possible extensions

- *pluginExtensionRule* : PLUGIN $\mapsto$ RULE
  returns the behavior of the plugin on extension points it is registered for.

- *pluginCallPriority* : PLUGIN $\times$ ENGINEMODE $\times$ ENGINEMODE $\mapsto$ NUMBER
  is the call priority of the plugin on the extension point between the two engine modes. Zero (0) is the lowest priority and 100 is the highest call priority. The engine will consider this priority when calling plugins at extension point transitions. Default call priority is 50.

The *Signature* and *IO* plugins from the standard CoreASM library, among others, implement the Extension Point interface to extend the control state ASM of the engine. We will look into these plugins in more detail in sections 6.4.1 and 6.4.3.

5.5.6 Plugin Service Interface

In many cases, there is a legitimate need for the environment of the CoreASM engine (e.g., the GUI of a simulator or of a debugger) to interact directly with some plugins. To support this interaction, the CoreASM extensibility framework introduces the concept of a *Plugin Service Interface* through which plugins can expose part of their functionality to the environment of the engine.

pluginServiceInterface : PLUGIN $\mapsto$ PLUGINSERVICEINTERFACE

The Plugin Service Interface allows CoreASM plugins to define and provide their own interfaces to the environment. Applications utilizing the engine can access these interfaces

through Control API and directly interact with such plugins. As an example, the IO Plugin provides its own interface to expose the output of its **print** rules to the environment of the engine (see Section 6.4.3). A GUI for the engine, for example, can utilize this interface to obtain the printed output and display it in a console window.

As each plugin exposes different functionalities, users of the Plugin Service Interface have to know in advance what to expect from a specific plugin. This requirement is in keeping with the assumption that the environment will access specific services from a specific plugin, as in the case of **print** rules.

5.5.7 Plugin Background

We model CoreASM plugins by elements of a domain PLUGIN. In addition to the special-purpose functions mentioned in this chapter, the following functions define a general interface for all plugins:

- *pluginName* : PLUGIN $\mapsto$ NAME
  returns the unique name of a plugin. The engine cannot load two plugins that share the same name.

- *pluginVersion* : PLUGIN $\mapsto$ VERSION
  returns the version information of the given plugin.

- *pluginDependencySet* : PLUGIN $\mapsto$ SET(NAME $\times$ VERSION)
  is a set of the names and minimum required version of all the plugins that this plugin depends on.

- *pluginLoadPriority* : PLUGIN $\mapsto$ NUMBER
  returns the suggested loading priority of this plugin. Zero (0) is the lowest priority and 100 is the highest loading priority. The engine will consider this priority when loading plugins. All plugins with the same priority level will be loaded in a non-deterministic order.

- *pluginInitRule* : PLUGIN $\mapsto$ RULE
  provides an ASM rule that initializes the plugin. This rule is called when the plugin is loaded by the engine; see the LoadPlugin rule on page 91.

For convenience, CoreASM allows plugins to be packaged together in one plugin, called a *package plugin*. For example, a set of standard CoreASM plugins (such as sets, numbers, and lists) can be packed in package plugin called the "Standard Plugin". If a plugin p is a package plugin, the value of *isPackagePlugin*(p) holds true and *enclosedPlugins*(p) returns the set of all the plugins enclosed in p.

Chapter 6

CoreASM: The Plugins

Most of the functionalities of CoreASM and its language constructs are provided through plugins to the CoreASM kernel. In this chapter we present the specification of those plugins that are currently available as part the CoreASM project. Most of these plugins are part of the standard library of CoreASM and can be loaded by simply loading the `Standard` package plugin.

Here, we divide the plugins into four categories: plugins that extend the CoreASM language by introducing new rule constructs (Section 6.1), plugins that provide the primitive data types such as numbers and character strings (Section 6.2), plugins that offer more complex data structures as collections of other elements (Section 6.3), and lastly, auxiliary plugins that extend the language and the engine with practically useful constructs and functionalities such as input/output mechanisms and scheduling policies (Section 6.4). The final section of this chapter introduces a special plugin, called JASMine, that allows access to Java objects and classes from CoreASM specifications.

Notation

Throughout this chapter, we use the pattern-action notation of Section 5.2.1 to formally define rule constructs, operators, and expression forms. In addition, we use the notation

```
foo: A -> B
```

in the description of a plugin p, denoting the extension of the vocabulary of the CoreASM state by plugin p through addition of a new Function element *fooFunction*, with the following

specification:

$$fooFunction \in \text{FUNCTIONELEMENT}$$
$$(\text{“foo”}, fooFunction) \in pluginFunctions(p)$$
$$signature(fooFunction) \equiv \langle \text{“A”}, \text{“B”} \rangle$$

6.1 Standard Rule Constructs

Abstract state machines come with a handful of standard control structures or transition rules (see Section 2.1.3). The most basic ASM rules (assignment, **import**, and **skip**) are defined in the kernel of the CoreASM engine as explained in Section 5.2.3. In this section, we extend the parser and the interpreter of the CoreASM engine through a number of rule plugins that provide the syntax and the semantics of standard and commonly-used ASM rule forms. The result of evaluating each rule, as we explained earlier, will be a multiset of update instructions that becomes the underlying value for the corresponding rule node in the parse tree.

We initiate by presenting rule plugins for all the rule forms defined for basic ASMs; we will then introduce plugins providing Turbo ASMs rule forms.

6.1.1 Block Rule Plugin

The most fundamental control structure in ASM is the block-rule, specified as follows:[1]

Block Rule

$$(\!| \{ {}^{\lambda_1}\boxed{r} \ \ldots {}^{\lambda_n}\boxed{r} \} |\!) \rightarrow \begin{array}{l} \textbf{choose } i \in [1..n] \textbf{ with } \neg evaluated(\lambda_i) \\ \quad pos := \lambda_i \\ \textbf{ifnone} \\ \quad [\![pos]\!] := (undef, \bigcup_{i \in [1..n]} updates(\lambda_i), undef) \end{array}$$

Here, all the rules in a block are evaluated in an unspecified order, with the final result being the multiset-union of all the update instructions produced by the various rules in the block.

[1] We provide here a rule for an n-elements block, whereas one for a two-elements block would suffice. Notice also that the same rule could be used for the alternative syntax R **par** Q, meaning that P and Q are to be executed in parallel. Finally, also note that we are disregarding here the scope constructors provided by the grammar—either relying on braces { } or on indentation to express nesting are common choices.

6.1.2 Conditional Rule Plugin

Close in importance comes the conditional rule construct, or the **if-then-else** rule. We accept a slightly extended syntax, where the guard is not restricted to be a *formula* (basically a Boolean predicate, as per Definition 2.4.14 in [25]), but rather any expression that may return **true**. This guarantees that plugins will be able to extend the set of allowable guards if needed. Notice that this approach is conservative with respect to the standard definition, given that formulae in the sense of [25] are indeed expressions supported by the Predicate Logic plugin (Section 6.2.1) in the CoreASM standard library.

Conditional Rule

$$(\!|\ \textbf{if}\ {}^{\alpha}\boxed{e}\ \ \textbf{then}\ {}^{\beta}\boxed{\ }\ |\!) \quad \rightarrow \quad pos := \alpha$$

$$(\!|\ \textbf{if}\ {}^{\alpha}v\ \textbf{then}\ {}^{\beta}\boxed{\ }\ |\!) \quad \rightarrow \quad \textbf{if}\ v = \mathsf{true}_e\ \textbf{then}\ pos := \beta\ \textbf{else}\ [\![pos]\!] := (undef, \{\!\}, undef)$$

$$(\!|\ \textbf{if}\ {}^{\alpha}v\ \textbf{then}\ {}^{\beta}u\ |\!) \quad \rightarrow \quad [\![pos]\!] := (undef, u, undef)$$

$$(\!|\ \textbf{if}\ {}^{\alpha}\boxed{e}\ \ \textbf{then}\ {}^{\beta}\boxed{\ }\ \textbf{else}\ {}^{\gamma}\boxed{\ }\ |\!) \quad \rightarrow \quad pos := \alpha$$

$$(\!|\ \textbf{if}\ {}^{\alpha}v\ \textbf{then}\ {}^{\beta}\boxed{\ }\ \textbf{else}\ {}^{\gamma}\boxed{\ }\ |\!) \quad \rightarrow \quad \textbf{if}\ v = \mathsf{true}_e\ \textbf{then}\ pos := \beta\ \textbf{else}\ pos := \gamma$$

$$(\!|\ \textbf{if}\ {}^{\alpha}v\ \textbf{then}\ {}^{\beta}u\ \textbf{else}\ {}^{\gamma}\boxed{\ }\ |\!) \quad \rightarrow \quad [\![pos]\!] := (undef, u, undef)$$

$$(\!|\ \textbf{if}\ {}^{\alpha}v\ \textbf{then}\ {}^{\beta}\boxed{\ }\ \textbf{else}\ {}^{\gamma}u\ |\!) \quad \rightarrow \quad [\![pos]\!] := (undef, u, undef)$$

6.1.3 The let-rule Plugin

The **let**-rule construct allows the definition of *environment* (read-only) variables (also called *logical* variables) which are not defined in the ASM state, but in a finite local environment. Once defined, the value of a logical variable cannot be updated by a transition rule.

Let Rule

$$(\!|\ \textbf{let}\ {}^{\alpha}x = {}^{\beta}\boxed{e}\ \ \textbf{in}\ {}^{\gamma}\boxed{\ }\ |\!) \quad \rightarrow \quad pos := \beta$$

$$(\!|\ \textbf{let}\ {}^{\alpha}x = {}^{\beta}v\ \textbf{in}\ {}^{\gamma}\boxed{\ }\ |\!) \quad \rightarrow \quad pos := \gamma$$
$$\mathsf{AddEnv}(x, v)$$

$$(\!|\ \textbf{let}\ {}^{\alpha}x = {}^{\beta}v\ \textbf{in}\ {}^{\gamma}u\ |\!) \quad \rightarrow \quad \mathsf{RemoveEnv}(x)$$
$$[\![pos]\!] := (undef, u, undef)$$

In a **let**-rule of the form '**let** $x = e$ **in** R' the scope of the logical variable x is the rule R but not the expression e.

6.1.4 The extend-rule Plugin

The **extend** rule is a syntactical sugar that imports a new element and adds it to a universe (extends the universe) [25, Table 2.4]. The semantics of an **extend**-rule of the form '**extend** U **with** x **do** R' is as follows: a new element is created and put in a logical variable x, the given rule R is evaluated, and the result of the evaluation of the **extend**-rule will be the union of the update multiset of its inner rule and a single update that adds the new element to universe U.

ExtendRule

$(\!\|$ **extend** $^{\alpha}\boxed{e}$ **with** $^{\beta}x$ **do**$^{\gamma}\boxed{r}$ $\|\!)$	$\rightarrow$	$pos := \alpha$

$(\!\|$ **extend** $^{\alpha}v$ **with** $^{\beta}x$ **do**$^{\gamma}\boxed{r}$ $\|\!)$	$\rightarrow$	**if** $isUniverse(v)$ **then**
		  $pos := \gamma$
		  let $e = new(\text{ELEMENT})$ **in**
		  AddEnv(x, e)
		else
		  Error('Extending a non-universe.')

$(\!\|$ **extend** $^{\alpha}v$ **with** $^{\beta}x$ **do**$^{\gamma}u$ $\|\!)$	$\rightarrow$	RemoveEnv(x) **let** $u' = u \cup \{\langle uniLoc(v, e), \mathsf{true}_e, updateAction\rangle\}$ **in** $[\![pos]\!] := (undef, u', undef)$

where

$uniLoc(v, e) \equiv (name, \langle e \rangle)$ s.t. $stateUniverse(state, name) = v$

6.1.5 The choose-rule Plugin

The **choose**-rule has the form '**choose** $x \in X$ **with** φ **do** R' where X is a collection of elements, φ is a Boolean expression and R is a rule. The semantics of the rule is execute R with an arbitrary element x from X that satisfies φ. In CoreASM, we extend this rule form by an optional **ifnone** clause that acts as an 'else' part: if no such element can be found the **ifnone** rule will be evaluated. We present here a simple form of **choose**-rule, with no additional condition on the chosen value and with an existing **ifnone** clause. A more comprehensive semantic definition is provided in Appendix A.5.1.

Choose Rule

$(\!\|$ **choose** $^{\alpha}x$ **in** $^{\beta}$[e] **do**$^{\gamma}$[r] **ifnone** $^{\delta}$[r] $\|\!)$	$\rightarrow$	$pos := \beta$
$(\!\|$ **choose** $^{\alpha}x$ **in** $^{\beta}v$ **do**$^{\gamma}$[r] **ifnone** $^{\delta}$[r] $\|\!)$	$\rightarrow$	**if** $enumerable(v)$ **then**
		let $s = enumerate(v)$ **in**
		if $\|s\| > 0$ **then**
		choose $t \in s$ **do**
		AddEnv(x, t)
		$pos := \gamma$
		else
		$pos := \delta$
		else
		Error('Choosing from a non-enumerable.')
$(\!\|$ **choose** $^{\alpha}x$ **in** $^{\beta}v$ **do**$^{\gamma}u$ **ifnone** $^{\delta}$[r] $\|\!)$	$\rightarrow$	RemoveEnv(x)
		$[\![pos]\!] := (undef, u, undef)$
$(\!\|$ **choose** $^{\alpha}x$ **in** $^{\beta}v$ **do**$^{\gamma}$[r] **ifnone** $^{\delta}u\|\!)$	$\rightarrow$	$[\![pos]\!] := (undef, u, undef)$

6.1.6 The forall-rule Plugin

The semantic definition of **forall**-rule is similar to that of **choose**-rule with the difference that all the elements of the given enumerable element that satisfy the optional guard are given a chance to be the free variable in the **do**-rule. Here, we present the semantics of **forall**-rule with a guard. The semantics of **forall** with no guard is presented in Appendix A.5.2.

Forall Rule

$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}$[e] $_1$ **with** $^{\gamma}$[e] $_2$ **do**$^{\delta}$[r] $\|\!)$	$\rightarrow$	$pos := \beta$
		$[\![pos]\!] := (undef, \{\!\}, undef)$
		$considered(\beta) := \{\}$
$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}v_1$ **with** $^{\gamma}$[e] $_2$ **do**$^{\delta}$[r] $\|\!)$	$\rightarrow$	**if** $enumerable(v_1)$ **then**
		let $s = enumerate(v_1) \backslash considered(\beta)$ **in**
		if $\|s\| > 0$ **then**
		choose $t \in s$ **do**
		AddEnv(x, t)
		$considered(\beta) := considered(\beta) \cup \{t\}$
		$pos := \gamma$
		else
		Error('Forall on a non-enumerable element')

$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}v_1$ **with** $^{\gamma}v_2$ **do**$^{\delta}$ [?] $\|\!)$	$\rightarrow$	**if** $v_2 = \mathsf{true}_e$ **then**
		$pos := \delta$
		else
		$pos := \beta$
		RemoveEnv(x)
		ClearTree(γ)

$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}v_1$ **with** $^{\gamma}v_2$ **do**$^{\delta}u$ $\|\!)$	$\rightarrow$	$pos := \beta$ RemoveEnv(x) ClearTree(γ) ClearTree(δ) $[\![pos]\!] := (undef, updates(pos) \cup u, undef)$

Notice that *considered* is used to keep track of values already considered for assignment to the free variable.

6.1.7 The case-rule Plugin

We present here the specification for a plugin implementing a parallel form of a switch case rule. The syntax is similar to the one that is used in [118],[2] but the semantics is quite different. Instead of evaluating the first rule with a matching guard value, all the rules with matching guard values will be evaluated in parallel. In essence, this parallel-case rule acts as a block rule in which all child rules are guarded against a given value.

To evaluate this rule, the case condition will be evaluated first and then all the guards will be evaluated in an unspecified order. Afterward, rules with a guard value equal to the value of the case condition will be evaluated. Finally, the updates generated by the matching cases are united to form the set of updates generated by the parallel-case rule. Formally, the construct is defined as follows:

Case Rule

$(\!\|$ **case** $^{\alpha}$[?] **of** $\{^{\lambda_1}$[?]$_1 : {}^{\lambda'_1}$[?]$_1 \ldots {}^{\lambda_n}$[?]$_n : {}^{\lambda'_n}$[?]$_n\}\|\!) \rightarrow \quad pos := \alpha$

$(\!\|$ **case** $^{\alpha}v$ **of** $\{^{\lambda_1}$[?]$_1 : {}^{\lambda'_1}$[?]$_1 \ldots {}^{\lambda_n}$[?]$_n : {}^{\lambda'_n}$[?]$_n\}\|\!) \rightarrow$
choose i **in** $[1..n]$ **with** $\neg evaluated(\lambda_i)$
$pos := \lambda_i$

[2]Here we use colons (:) instead of arrows ($\rightarrow$).

$$(\!|\textbf{case } {}^{\alpha}v \textbf{ of } \{{}^{\lambda_1}v_1 : {}^{\lambda'_1}\boxed{?}_1 \ldots {}^{\lambda_n}v_n : {}^{\lambda'_n}\boxed{?}_n\}|\!) \rightarrow$$
$$\textbf{choose } i \textbf{ in } [1..n] \textbf{ with } equal(v, v_i) \wedge \neg evaluated(\lambda'_i)$$
$$pos := \lambda'_i$$
$$\textbf{ifnone}$$
$$[\![pos]\!] := (undef, \bigcup_{i \in [1..n] \wedge equal(v,v_i)} updates(\lambda'_i), undef)$$

6.1.8 The TurboASM Plugin

Basic ASMs are further extended by operators for sequential composition and iteration of ASMs, and also by parameterized submachines [25]. These extended ASMs are called *Turbo ASMs*. Following the definitions of those operators, the TurboASM plugin provides sequentiality and iteration rule forms, together with support for local state definitions and constructs allowing rules to return values.

The seq-rule

Sequential composition of rules is facilitated by the **seq**-rule acting as an operator on rules. According to [25, Def. 4.1.1], the semantics of 'P **seq** Q' is defined as the effect of first executing P in the current state $\mathfrak{A}$, and then executing Q in the resulting state $\mathfrak{A} + U_P$ where U_P is the update set produced by P. If U_P is inconsistent, the result of the sequence composition will be U_P.

Since we want to model the effect of evaluating the second rule in a sequence in the state that would be produced by applying the updates produced by the first rule, we have to "simulate" the application of the updates, without really modifying the current state. This is obtained by using a *stack* of states, managed through three macros: PushState copies the current state in the stack, PopState retrieves the state from the top of the stack (thus discarding the current state), and Apply(u) applies the updates in the update set u to the current state. Formal definitions for these macros are given in Appendix A.1. Based on the intuitive understanding of these macros, the interpreter plugin for the **seq**-rule can be specified as follows:

SeqRule

$(\!\vert {}^{\alpha}\boxed{r}_1 \textbf{ seq } {}^{\beta}\boxed{r}_2 \vert\!)$	$\rightarrow$	$pos := \alpha$

$(\!\vert {}^{\alpha}u_1 \textbf{ seq } {}^{\beta}\boxed{r}_2 \vert\!)$	$\rightarrow$	**let** $uSet$ = Aggregate(u_1) **in**
		if $isConsistent(uSet) \wedge aggregationConsistent(u_1)$ **then**
		PushState
		Apply$(uSet)$
		$pos := \beta$
		else
		$[\![pos]\!] := (undef, u_1, undef)$

$(\!\vert {}^{\alpha}u_1 \textbf{ seq } {}^{\beta}u_2 \vert\!)$	$\rightarrow$	**local** $uMset$ $[uMset \leftarrow$ Compose$(u_1, u_2)]$ **in** PopState $[\![pos]\!] := (undef, uMset, undef)$

Before consistency of the update instructions produced by the first rule can be checked, the resultant update instructions must be aggregated into regular updates. If both aggregation consistency and update set consistency hold, the resultant update set is applied to the current state producing a temporary state; otherwise the inconsistent update multiset is returned. If the update instructions produced by the first rule are consistent, the second rule is fired in the temporary state, resulting in the second update multiset. The first and second update multisets must then be sequentially composed. The update multiset resulting from the sequential composition is the update multiset produced by the **seq**-rule in the simulated machine.

In order to improve the readability of specifications, CoreASM provides the following syntax for the sequential composition of rules, in which the **next** keyword is optional:

$$\textbf{seq } P \textbf{ next } Q \;\equiv\; P \textbf{ seq } Q$$

The iterate Rule

The **iterate**-rule repeatedly executes its body, until the update set produced is either empty or inconsistent; at that point, the accumulated updates are computed. The resulting update set can be inconsistent if the computation of the last step had produced an inconsistent set of updates. The semantic definition is similar in principle to that of the **seq**-rule:

Iterate Rule

$(\!|\ \textbf{iterate}\ ^{\alpha}\boxed{r}\ |\!) \rightarrow$ PushState
$composedUpdates(pos) := \{\!\}$
$pos := \alpha$

$(\!|\ \textbf{iterate}\ ^{\alpha}u\ |\!) \rightarrow$ **if** $u \neq \{\!\}$ **then**
let $uSet$ = Aggregate(u),
$composed \leftarrow$ Compose$(composedUpdates(pos), u)$ **in**
$composedUpdates(pos) := composed$
if $aggregationConsistent(u) \wedge isConsistent(uSet)$ **then**
Apply$(uSet)$
ClearTree(α)
$pos := \alpha$
else
PopState
$[\![pos]\!] := (undef, composed, undef)$
else
PopState
$[\![pos]\!] := (undef, composedUpdates(pos), undef)$

Notice here how iteration is carried on in a separate state, after saving the original one in the stack. After the iteration is completed, the update instruction multisets are composed into a single multiset of update instructions to be applied to the initial state. The initial state is then restored from the stack, and the computed updates are assigned to the node. Also, notice that after each step in the iteration, the entire subtree is cleared (i.e., the $[\![\cdot]\!]$ function of each node is set to *undef*), so that the computation of the next step can proceed on a clean parse tree.

The while Rule

The non-standard **while**-rule can also be defined in a similar way. The semantics of a rule '**while** (*cond*) R' is to iterate the execution of R as long as *cond* evaluates to true and R does not produce an empty or inconsistent update set. Thus, the following equivalence holds:

$$\textbf{while}\ (cond)\ R\ \equiv\ \textbf{iterate if}\ cond\ \textbf{then}\ R$$

Thus, the semantics of the **while** rule closely follows that of the **iterate** rule:

While Rule

$(\!|\ \textbf{while}\ ({}^{\alpha}\boxed{c}\ )\ {}^{\beta}r\ |\!) \rightarrow$

PushState
$composedUpdates(pos) := \{\!|\}$
$pos := \alpha$

$(\!|\ \textbf{while}\ ({}^{\alpha}v)\ {}^{\beta}\boxed{r}\ |\!) \rightarrow$

if $v = \mathsf{true}_e$ **then**
  $pos := \beta$
else
  PopState
  $[\![pos]\!] := (undef, composedUpdates(pos), undef)$

$(\!|\ \textbf{while}\ ({}^{\alpha}v)\ {}^{\beta}u\ |\!) \rightarrow$

if $u \neq \{\!|\}$ **then**
  let $uSet$ = Aggregate(u),
    $composed \leftarrow$ Compose$(composedUpdates(pos), u)$ **in**
  $composedUpdates(pos) := composed$
  if $aggregationConsistent(u) \wedge isConsistent(uSet)$ **then**
    Apply$(uSet)$
    ClearTree(α)
    ClearTree(β)
    $pos := \alpha$
  else
    PopState
    $[\![pos]\!] := (undef, composed, undef)$
else
  PopState
  $[\![pos]\!] := (undef, composedUpdates(pos), undef)$

Notice that other choices for the semantics of **while** were also possible: for example, [25, Example 4.1.4] presents a variant that does not terminate when the update set produced by the rule is empty (their Example 4.1.2 is instead consistent with our definition).

More generally, both **iterate** and **while** could also be defined to terminate when the update set contributed by the body of the rule does not modify the state. To our knowledge, this semantics has not been explored and applied in practice.

Local State and Return Values

Local state is introduced in rules by a special syntax [25, Def. 4.1.5] which introduces local state function names together with their initialization rules. Updates made to these special

locations are then discarded before returning the final update set to the caller. In the same spirit, return values are simulated by designating a special location in the state, and by using the last update to that location as return value.

We sketch here only the basic idea of how local state and return values are handled. In particular, we omit the details of how local state initialization is performed, based on the observation that a declaration of local state with initialization can be transformed into a declaration without initialization followed by an explicit sequential composition of an assignment and the main rule.

Local Rule

$(\!|\, \textbf{local}\ {}^{\lambda_1} x_1 \ldots {}^{\lambda_n} x_n\ \textbf{in}\ {}^{\alpha}\boxed{l}\ |\!) \quad \rightarrow \quad pos := \alpha$

$(\!|\, \textbf{local}\ {}^{\lambda_1} x_1 \ldots {}^{\lambda_n} x_n\ \textbf{in}\ {}^{\alpha} u\, |\!) \quad \rightarrow \quad [\![pos]\!] := (undef, u \ominus \{x_1, \ldots, x_n\}, value(\alpha))$

---

where the $\ominus$ operator is defined as follows:

$$U \ominus H = \{\langle l, v, a\rangle \in U \mid name_{lc}(l) \notin H\}$$

A frequent and idiomatic use of Turbo ASMs is to compute functions by executing a rule and then extracting a value from the resulting set of updates, rather than applying the updates to the state. The semantics of the following Turbo ASM call with return values

$$l \leftarrow \mathsf{R}(a_1, \ldots, a_n)$$

is to replace every occurrence of a special variable **result** in the body of the rule R with l, and call rule R [25, Def. 4.1.7]. The following pattern provides a formal semantics for this rule form in CoreASM:

Return Result Rule

$(\!|\, {}^{\alpha}\boxed{l} \leftarrow {}^{\beta} x({}^{\lambda_1}\boxed{?}_1, \ldots, {}^{\lambda_n}\boxed{?}_n)\, |\!) \quad \rightarrow \quad \textbf{if}\ isRuleName(x)\ \textbf{then}$
$\mathsf{ReturnResultRuleCall}(ruleValue(x), \langle \lambda_1, \ldots, \lambda_n\rangle, l)$

The ReturnResultRuleCall routine, defined below, describes how calls to rules with the special **result** location are handled in CoreASM.

Turbo ASM Plugin

ReturnResultRuleCall$(r, args, l) \equiv$
  if $workCopy(pos) = undef$ **then**
    let $params = concat(\text{“result”}, param(r)), args = concat(l, args)$ **in**
      let $b' = \mathsf{CopyTreeSub}(body(r), param(r), args)$ **in**
        $workCopy(pos) := b'$
        $parent(b') := pos$
        $pos := b'$
  else
    $[\![pos]\!] := (undef, updates(workCopy(pos)), value(workCopy(pos)))$
    $workCopy(pos) := undef$

The syntax provided above, however, is not particularly practical, as the computation is restricted to be a statement assigning a value to a given identifier, and so cannot be used inside a complex expression. For example, one has to write

$x \leftarrow R(a_1, \ldots, a_n)$
$y \leftarrow Q(b_1, \ldots, b_m)$
seq
$z := x + y$

instead of the more natural

$z := R(a_1, \ldots, a_n) + Q(b_1, \ldots, b_m)$

Hence, we propose here an alternative syntax and semantics of the form

return e **in** R

in which e is an expression and R is a rule. The semantics of this construct is to execute R in the current state $\mathfrak{A}$ and if the resulting update multiset is consistent, evaluate e in the state $\mathfrak{A}+U_R$ (where U_R is the updates produced by R) and return the value of e, discarding U_R. We formally describe this semantics in the following rules:

ReturnRule

$(\!\|$ **return** $^{\alpha}\boxed{e}$ **in** $^{\beta}r \|\!)$	$\rightarrow$	$pos := \beta$

$(\!\|$ **return** $^{\alpha}\boxed{e}$ **in** $^{\beta}u \|\!)$	$\rightarrow$	**let** $uSet$ = Aggregate(u) **in**
		if $isConsistent(uSet) \wedge aggregationConsistent(u)$ **then**
		PushState
		Apply$(uSet)$
		$pos := \alpha$
		else
		$[\![pos]\!] := (undef, \{\!\}, \mathsf{undef}_e)$

$(\!\|$ **return** $^{\alpha}v$ **in** $^{\beta}u \|\!)$	$\rightarrow$	PopState $[\![pos]\!] := (undef, \{\!\}, v)$

In this construct, the rule r is executed first; the return expression is evaluated in the state obtained by provisionally applying the updates from r to the current state, and the resulting value is returned, while the updates and the provisional state itself are discarded.

6.2 Primitive Data Types

In this section we introduce those plugins that extend the CoreASM engine with backgrounds of primitive data types, basically numbers and character strings. We also include in this section the Predicate Logic plugin that offers Boolean operators defined on Boolean elements introduced in the CoreASM kernel.

6.2.1 The Predicate Logic Plugin

The Predicate Logic plugin provides operators implementing a Boolean algebra. Since the corresponding background is already provided by the kernel, this plugin extends only the parser and the interpreter of the CoreASM engine to provide the standard Boolean operators together with the universal and the existential quantifiers.

The only unary operator provided by this plugin is the negation operator: **not**. The semantics of this operator is very simple and is formally defined by the following rule:

Predicate Logic Plugin: not

$$\langle\!|\ \textbf{not}\ {}^{\alpha}\boxed{?}\ |\!\rangle_{[850]} \rightarrow \begin{array}{l} \textbf{if}\ \neg \mathit{evaluated}(\alpha)\ \textbf{then} \\ \quad \mathit{pos} := \alpha \\ \textbf{else} \\ \quad \textbf{if}\ \mathit{isBoolean}(\mathit{value}(\alpha))\ \textbf{then} \\ \quad\quad \textbf{if}\ \mathit{value}(\alpha) = \mathsf{true}_e\ \textbf{then} \\ \quad\quad\quad [\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathsf{false}_e) \\ \quad\quad \textbf{else} \\ \quad\quad\quad [\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathsf{true}_e) \end{array}$$

The Predicate Logic plugin also provides the standard binary operators **and**, **or**, **xor**, and **implies**, together with the not-equality operator `!=`. As an example, we present here the semantic definition of the logical implication operator:

Predicate Logic Plugin: implies

$$\langle\!|\ {}^{\alpha}\boxed{?}\ \textbf{implies}\ {}^{\beta}\boxed{?}\ |\!\rangle_{[375]} \rightarrow \begin{array}{l} \textbf{choose}\ \lambda \in \{\alpha, \beta\}\ \textbf{with}\ \neg \mathit{evaluated}(\lambda) \\ \quad \mathit{pos} := \lambda \\ \textbf{ifnone} \\ \quad \textbf{if}\ \mathit{isBoolean}(\mathit{value}(\alpha)) \wedge \mathit{isBoolean}(\mathit{value}(\beta))\ \textbf{then} \\ \quad\quad \textbf{if}\ ((\mathit{value}(\alpha) = \mathsf{false}_e) \vee (\mathit{value}(\beta) = \mathsf{true}_e))\ \textbf{then} \\ \quad\quad\quad [\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathsf{true}_e) \\ \quad\quad \textbf{else} \\ \quad\quad\quad [\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathsf{false}_e) \end{array}$$

In addition, the Predicate Logic plugin also provides the membership operator '$\in$'. If the operand on the right hand side (rhs) is an enumerable, this operator returns true if that enumerable includes the operand on the left hand side (lhs). We have:

Predicate Logic Plugin: memberof

$$\langle\!|\ {}^{\alpha}\boxed{?}\ \textbf{memberof}\ {}^{\beta}\boxed{?}\ |\!\rangle_{[550]} \rightarrow \begin{array}{l} \textbf{choose}\ \lambda \in \{\alpha, \beta\}\ \textbf{with}\ \neg \mathit{evaluated}(\lambda) \\ \quad \mathit{pos} := \lambda \\ \textbf{ifnone} \\ \quad \textbf{if}\ \mathit{enumerable}(\mathit{value}(\alpha))\ \textbf{then} \\ \quad\quad \textbf{if}\ \mathit{value}(\beta) \in \mathit{enumerate}(\mathit{value}(\alpha))\ \textbf{then} \\ \quad\quad\quad [\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathsf{true}_e) \\ \quad\quad \textbf{else} \\ \quad\quad\quad [\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathsf{false}_e) \end{array}$$

The formal definition of other operators is available in Appendix A.5.3.

Two logical quantifiers $\exists$ and $\forall$ are also provided by the Predicate Logic plugin with the following syntax

$$\textbf{exists } x \textbf{ in } X \textbf{ with } \varphi$$
$$\textbf{forall } x \textbf{ in } X \textbf{ holds } \varphi$$

in which X is an enumerable, φ is a Boolean predicate and the scope of x is limited to φ. We present here the semantic definition of the existential quantifier. The definition of the universal quantifier is very similar and is presented in Appendix A.5.3. Notice again the use of the *considered* function to keep track of the elements that we considered so far.

Predicate Logic Plugin: exists

$(\!|\ \textbf{exists}^{\alpha} x \textbf{ in } {}^{\beta}\boxed{e} \textbf{ with } {}^{\gamma}\boxed{e}\ |\!) \rightarrow$

$pos := \beta$

$considered(\beta) := \{\}$

$(\!|\ \textbf{exists}^{\alpha} x \textbf{ in } {}^{\beta}v \textbf{ with } {}^{\gamma}\boxed{e}\ |\!) \rightarrow$

```
if enumerable(v) then
  let s = enumerate(v)\considered(β) in
    if |s| > 0 then
      choose t ∈ s do
        AddEnv(x, t)
        considered(β) := considered(β) ∪ {t}
      pos := γ
    else
      [[pos]] := (undef, undef, false_e)
else
  Error('Cannot enumerate a non-enumerable element')
```

$(\!|\ \textbf{exists}^{\alpha} x \textbf{ in } {}^{\beta}v \textbf{ with } {}^{\gamma}v\ |\!) \rightarrow$

```
if (value(γ) = true_e) then
  [[pos]] := (undef, undef, true_e)
else
  pos := β
RemoveEnv(x)
ClearTree(γ)
```

6.2.2 The Number Plugin

The Number plugin extends the abstract storage, the parser, and the interpreter of the CoreASM engine to provide the *Number* background, representing the domain of Real numbers $\mathbb{R}$, together with necessary functions and operators needed to work with both integer and real numbers. The background of Number elements is defined as $numberBkg \in$ BACKGROUNDELEMENT; we have

$$name(numberBkg) = \text{“NUMBER”}$$
$$newValue(numberBkg) = zero$$

Number elements are values of the domain NUMBERELEMENT. We have

$$\forall ne \in \text{NUMBERELEMENT} \quad member_{ue}(numberBkg, n) = true$$

We define the following functions to provide a mapping from Number elements to the actual numeric values they represent and vice versa:

$$numberElement : \mathbb{R} \mapsto \text{NUMBERELEMENT}$$
$$numericValue : \text{NUMBERELEMENT} \mapsto \mathbb{R}$$

Finally, the equality of two Number elements is defined as the equality of the numeric values they represent (see also Section 5.1):

$$\forall ne' \in \text{NUMBERELEMENT} \quad equal_{Number}(ne, ne') \equiv numericValue(ne) = numericValue(ne')$$

Operators

The Number plugin provides the following numeric operators:

- “+” : the addition binary operator (precedence level: 750)
- “-” : the subtraction binary operator (precedence level: 750)
- “-” : the negation unary operator (precedence level: 850)
- “*” : the multiplication binary operator (precedence level: 800)
- “/” : the division binary operator (precedence level: 800)
- “div” : the integer division binary operator (precedence level: 800)
  a **div** $b \equiv floor(a/b)$

- "%" : the modulus (remainder) binary operator (precedence level: 800)
  $a \ \% \ b \equiv \mathit{floor}(a/b)$
- "^" : the exponential binary operator (precedence level: 820)

We present here the semantics of the addition operator (i.e., "+"). The same approach is used to define the rest of the above operators.

Number Plugin

$(\!|\ ^{\alpha}[?] + ^{\beta}[?]\ |\!)_{[750]}$ →

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg \mathit{evaluated}(\lambda)$
  $\mathit{pos} := \lambda$
ifnone
  if $l \in$ NUMBERELEMENT $\wedge\ r \in$ NUMBERELEMENT **then**
    $[\![\mathit{pos}]\!] := (\mathit{undef}, \mathit{undef}, \mathit{result})$
where
  $\mathit{result} \equiv \mathit{numberElement}(\mathit{numericValue}(l) + \mathit{numericValue}(r))$
  $l \equiv \mathit{value}(\alpha)$
  $r \equiv \mathit{value}(\beta)$

The Number plugin also provides the following relational operators defined on Number elements:

- ">" : greater-than binary operator (precedence level: 650)
- ">=" : greater-than or equal-to binary operator (precedence level: 650)
- "<" : less-than binary operator (precedence level: 650)
- "<=" : less-than or equal-to binary operator (precedence level: 650)

The greater-than operator is defined as follows:

Number Plugin

$(\!|\ {}^{\alpha}\boxed{?} > {}^{\beta}\boxed{?}\ |\!)_{[650]}$ →

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if $l \in$ NUMBERELEMENT $\wedge\ r \in$ NUMBERELEMENT **then**
    $[\![pos]\!] := (undef, undef, result)$
  else
    Error('Both operands must be numbers.')
where
  $result \equiv booleanValue(numericValue(l) > numericValue(r))$
  $l \equiv value(\alpha)$
  $r \equiv value(\beta)$

The semantics of the other three relational operators are also defined in a similar fashion.

Functions

The Number plugin extends the vocabulary of the state with the following two functions:

- `infinity: -> NUMBER`
  returns the positive infinity.
- `toNumber: ELEMENT -> NUMBER`
  if possible, maps the given element to a Number element it represents.

Number Classes

The Number plugin provides the user with the following predicates in order to identify whether a number belongs to a particular numerical class:

- `isNaturalNumber: NUMBER -> BOOLEAN`
  $fGetValue(isNaturalNumberFunction, \langle n \rangle) = \begin{cases} \mathsf{true}_e, & \text{if } numericValue(n) \in \mathbb{N}; \\ \mathsf{false}_e, & \text{otherwise.} \end{cases}$.
- `isIntegerNumber: NUMBER -> BOOLEAN`
  $fGetValue(isIntegerNumberFunction, \langle n \rangle) = \begin{cases} \mathsf{true}_e, & \text{if } numericValue(n) \in \mathbb{Z}; \\ \mathsf{false}_e, & \text{otherwise.} \end{cases}$.

- `isRealNumber: NUMBER -> BOOLEAN`
  $fGetValue(isRealNumberFunction, \langle n \rangle) = \begin{cases} \mathsf{true}_e, & \text{if } numericValue(n) \in \mathbb{R}; \\ \mathsf{false}_e, & \text{otherwise.} \end{cases}$.

Number Characteristics

To identify the characteristics of numbers, the following predicates are defined on all Number elements:

- `isEvenNumber: NUMBER -> BOOLEAN`
  $fGetValue(isEvenNumberFunction, \langle n \rangle) =$
  $\begin{cases} \mathsf{true}_e, & \text{if } numericValue(n) \in \mathbb{Z} \ \wedge \ numericValue(n)\%2 = 0; \\ \mathsf{false}_e, & \text{otherwise.} \end{cases}$.
- `isOddNumber: NUMBER -> BOOLEAN`
  $fGetValue(isOddNumberFunction, \langle n \rangle) =$
  $\begin{cases} \mathsf{true}_e, & \text{if } numericValue(n) \in \mathbb{Z} \ \wedge \ numericValue(n)\%2 = 1; \\ \mathsf{false}_e, & \text{otherwise.} \end{cases}$.

Number Ranges

Number plugin also provides the NumberRange background which is the background of number ranges of the form $[a..b : s]$ where a and b are respectively the starting and the ending values of the range (inclusive) and s is the *step* of the range. The background of Number Range elements is provided by $numberRangeBkg \in$ BackgroundElement, where

$$name(numberRangeBkg) = \text{“NUMBER\_RANGE”}$$
$$newValue(numberRangeBkg) = [0..1 : 1]$$

The following functions are defined on Number Range elements (see Section 5.1):

- $bkg(r) =$ “NumberRange” where $r \in$ NumberRange.
- $rangeFrom$: NumberRange $\mapsto$ Number
  holds the lower boundary of the Number Range element.
- $rangeTo$: NumberRange $\mapsto$ Number
  holds the upper boundary of the Number Range element.

- $rangeStep$: NumberRange $\mapsto$ Number
  holds the range step.

- $\forall nr_1, nr_2 \in$ NumberRange $\ equal_{NumberRange}(nr_1, nr_2) \equiv$
  $rangeFrom(nr_1) = rangeFrom(nr_2)$
  $\wedge\ rangeTo(nr_1) = rangeTo(nr_2)$
  $\wedge\ rangeStep(nr_1) = rangeStep(nr_2)$

- $\forall r \in$ NumberRange, $\ enumerable(r)$
  All Number Range elements are enumerable.

- $enumerate_{IntegerRange}$: NumberRange $\mapsto$ List(Element)
  provides a collection of Elements representing the numbers that are included in the given Number Range.

  $enumerate(r) \equiv [\ x \mid x = rangeFrom(r) + i * rangeStep(r)\ \wedge\ i \in \mathbb{N}\ \wedge\ x \leq rangeTo(r)]$

The following expression form creates a Number Range element:

Integer Range

$(\!|\ [^{\alpha}\square\ ..^{\beta}\square\ :^{\gamma}\square]\ |\!)$ $\rightarrow$

choose $\lambda \in \{\alpha, \beta, \gamma\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if $\forall v \in \{l, r, s\}\ \ isNumber(v)$ **then**
    let $newRange = newValue(numberRangeBack)$ **in**
      $rangeFrom(newRange) := numericValue(l)$
      $rangeTo(newRange) := numericValue(r)$
      $rangeStep(newRange) := numericValue(s)$
      $[\![pos]\!] := (undef, undef, newRange)$
  else
    Error('Both operands must be numbers.')
where
  $l \equiv value(\alpha)$
  $r \equiv value(\beta)$
  $s \equiv value(\gamma)$

In the above form, the step of a range (γ) can be omitted in which case it would be considered to be 1.

6.2.3 The String Plugin

The String plugin provides all that is needed to work with character strings as elements of the CoreASM state. The background of String elements is provided by $stringBack \in$ BACKGROUNDELEMENT; we have

$$name(stringBack) = \text{“STRING”}$$
$$newValue(stringBack) = emptyString$$

We model String elements as values of a domain STRINGELEMENT. The following functions are defined on String elements:

- $stringValue$: STRINGELEMENT $\mapsto$ LIST(CHARACTER)
  for every String element returns the sequence of characters in that string.

- $stringElement$: ELEMENT $\mapsto$ STRINGELEMENT
  maps every element to a String representation of that element. The exact semantics of this function depends on the Element itself and it is left abstract here.

- $concatString$: STRINGELEMENT $\times$ STRINGELEMENT $\mapsto$ STRINGELEMENT
  concatenates two string elements into one. For all $s_1, s_2 \in$ STRINGELEMENT, we have

$$concatString(s_1, s_2) \equiv concat(stringValue(s_1), stringValue(s_2))$$

For every $s \in$ STRINGELEMENT we have (see Section 5.1):

- $bkg(s) =$ “StringElement”
- $\forall s' \in$ STRINGELEMENT $\quad equal_{String}(s, s') \equiv stringValue(s) = stringValue(s')$
- $\forall s \in$ STRINGELEMENT, $\quad enumerable(s)$
  All String elements are enumerable.
- $enumerate_{String}(s) = l \in$ LIST(STRINGELEMENT)
  where l is a list of String elements representing the characters of s.

Operators

The String plugin provides the following concatenation operator on String elements:

String Plugin

$(\!|\ {}^{\alpha}\boxed{?} + {}^{\beta}\boxed{?}\ |\!)_{[750]}$ →

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if $l \in$ StringElement $\wedge\ r \in$ StringElement **then**
    $[\![pos]\!] := (undef, undef, concatString(l, r))$
where
  $l \equiv value(\alpha)$
  $r \equiv value(\beta)$

Functions

The String plugin extends the CoreASM state with the following two functions defined on String elements:

- `toString: ELEMENT -> STRING`
  returns a string representation of the given element. We have,
  $\forall e \in$ Element $\ value_{fe}(toStringFunction, \langle e \rangle) = stringElement(e)$
- `strlen: STRING -> NUMBER`
  returns the length of the given string. For all $s \in$ StringElement we have,
  $value_{fe}(strlenFunction, \langle s \rangle) = numberElement(|stringValue(s)|)$

The String plugin relies on the availability of the Number background provided by the Number plugin.

6.3 Collections

We use the term *collection* to refer to the most abstract concept of a grouping of zero or more elements with potential multiplicities of more than one. In this section, we introduce those CoreASM plugins that offer backgrounds implementing different kinds of collections. The most liberal implementation of collections in CoreASM is provided by the Bag plugin (Section 6.3.3). Other plugins, such as the Set plugin (Section 6.3.2) and the List plugin

(Section 6.3.4), offer more specialized forms of collections. The Collection plugin, introduced in Section 6.3.1, provides the foundation for collection backgrounds in CoreASM.

6.3.1 The Collection Plugin

The Collection plugin provides a cornerstone for collections in CoreASM, offering a set of common functions and rule forms defined on collections. However, each specific collection background (e.g., list or set) is provided separately by its corresponding plugin.

Modifiable Collections

The Collection plugin introduces a modifiable-collection attribute on elements, defined by the following function:

$$isModifiableCollection : \textsc{Element} \mapsto \textsc{Boolean}$$

The modifiability attribute set on an element indicates that generic collection modifications (at this point limited to addition and removal of an element) can be applied to the element. Plugins that provide modifiable collection elements (such as sets and list) must also provide the semantics of such modifications through two functions of the form

$$computeAddUpdate_{bkg} : \textsc{Location} \times \textsc{Element} \mapsto \textsc{Multiset}(\textsc{Update})$$
$$computeRemoveUpdate_{bkg} : \textsc{Location} \times \textsc{Element} \mapsto \textsc{Multiset}(\textsc{Update})$$

where *bkg* is the collection background the plugin provides. These two functions are expected to produce proper update instructions to add/remove elements to/from locations holding collection elements.

Rule Forms

The Collection plugin extends the CoreASM language with two rule forms for adding and removing elements to and from collections. As explained above, the semantics of these rule forms relies on the add and remove semantics provided by the plugin of each collection element.

Collection Plugin: Add-To

$(\!|\ \textbf{add}\ {}^{\alpha}\boxed{e}\ \ \textbf{to}\ {}^{\beta}\boxed{l}\ |\!) \rightarrow$

choose $\tau \in \{\alpha, \beta\}$ **with** $\neg evaluated(\tau)$
  $pos := \tau$
ifnone
  let $c = value(\beta)$ **in**
    if $isModifiableCollection(c)$ **then**
      let $u = computeAddUpdate_{bkg(c)}(loc(\beta), value(\alpha))$ **in**
        $[\![pos]\!] := (undef, u, undef)$

Collection Plugin: Remove-From

$(\!|\ \textbf{remove}\ {}^{\alpha}\boxed{e}\ \ \textbf{from}\ {}^{\beta}\boxed{l}\ |\!) \rightarrow$

choose $\tau \in \{\alpha, \beta\}$ **with** $\neg evaluated(\tau)$
  $pos := \tau$
ifnone
  let $c = value(\beta)$ **in**
    if $isModifiableCollection(c)$ **then**
      let $u = computeRemoveUpdate_{bkg(c)}(loc(\beta), value(\alpha))$ **in**
        $[\![pos]\!] := (undef, u, undef)$

Functions

The Collection plugin also provides the following functions defined on enumerable elements:

- `foldl: ELEMENT * FUNCTION * ELEMENT -> ELEMENT`
  which implements the following function:
  $foldl([x_1, \ldots, x_n], f, i) \equiv f(x_n, f(x_{n-1}, \ldots f(x_1, i))) \ldots)$
- `foldr: ELEMENT * FUNCTION * ELEMENT -> ELEMENT`
  which implements the following function:
  $foldr([x_1, \ldots, x_n], f, i) \equiv f(x_1, f(x_2, \ldots f(x_n, i))) \ldots)$
- `fold: ELEMENT * FUNCTION * ELEMENT -> ELEMENT`
  is the same as `foldr`.
- `fold: ELEMENT * FUNCTION -> ELEMENT`
  which implements the following function:
  $map([x_1, \ldots, x_n], f) \equiv [f(x_1), f(x_2), \ldots f(x_n)]$

- `filter: ELEMENT * FUNCTION -> ELEMENT`
  which implements the following function:
  $filter(\{x_1, \ldots, x_n\}, f) \equiv \{x_i \mid f(x_i)\}$
  $filter([x_1, \ldots, x_n], f) \equiv [x_i \mid f(x_i)]$

The Collection plugin depends on the availability of the Number background provided by the Number plugin.

6.3.2 The Set Plugin

The Set plugin extends the CoreASM state by providing the background of sets with its operations and functions.[3] The background of Set elements is provided by $setBack \in$ BACKGROUNDELEMENT; we have

$$name(setBack) = \text{“SET”}$$
$$newValue(setBack) = emptySet$$

Set elements are values of the domain SETELEMENT. The following functions define the interface of Set elements by providing a mapping between Set elements and the actual set of elements they represent:

- $setElement$: SET(ELEMENT) $\mapsto$ SETELEMENT
  for every set of elements, returns a Set element representation of that set.

- $setMembers$: SETELEMENT $\mapsto$ SET(ELEMENT)
  for every Set element, returns the set of its members.

For all $s \in$ SETELEMENT we have:

- $bkg(s) :=$ “Set”

- $\forall s' \in$ SETELEMENT $\quad equal_{Set}(s, s') \equiv setMembers(s) = setMembers(s')$

- $enumerable(s)$
  All Set elements are enumerable.

[3] This section is based on Mashaal Memon's M.Sc. work previously published in [99] with improvements and modifications.

- $enumerate_{Set}(s) = setMembers(s)$.
- $s \in$ FUNCTIONELEMENT
  All Set elements also behave as functions.
- $class_{fe}(s) = static$
- $\forall e \in$ ELEMENT $\quad value_{fe}(s, \langle e \rangle) \equiv booleanValue(e \in setMembers(s))$

To facilitate partial updates to sets, the **add/to**-rule and **remove/from**-rule are supported by the Set plugin (see Section 6.3.1). We have

$$\forall s \in \text{SETELEMENT}\ isModifiableCollection(s)$$

The single addition of an element from a set, or the **add/to**-rule, results in an instruction to carry out a *setAddAction* action; the removal of a single element from a set, or the **remove/from**-rule, results in an instruction to perform a *setRemoveAction* action. For all $loc \in Location$ and $value \in Element$, we have

$$computeAddUpdate_{Set}(loc, value) \equiv \{\!|\langle loc, value, setAddAction\rangle|\!\}$$
$$computeRemoveUpdate_{Set}(loc, value) \equiv \{\!|\langle loc, value, setRemoveAction\rangle|\!\}$$

Notice that no checks are made to ensure that the value of the location is in fact a set. This is deferred to the aggregation phase.

Set Enumeration and Comprehension

The set plugin provides two methods of set description: namely set enumeration and set comprehension. With the former, one is able to explicitly describe the contents of a set by listing its individual elements:

Set Plugin: Set Enumeration

$(\!|\{\ {}^{\lambda_1}\boxed{?}_1, \ldots, {}^{\lambda_n}\boxed{?}_n\ \}|\!) \quad \rightarrow$

choose $i \in [1..n]$ **with** $\neg evaluated(\lambda_i)$
  $pos := \lambda_i$
ifnone
  let $s = \{value(\lambda_i) \mid i \in [1..n]\}$ **in**
    $[\![pos]\!] := (undef, undef, setElement(s))$

The latter allows one to describe set contents algorithmically. There are many accepted syntactic and semantic variants; the Set plugin provides three variants which we believe

encompass a wide range of algorithmically expressible finite sets. Given a set comprehension expression of the form

$$\{x_0 \textbf{ is } exp_0 \mid x_1 \textbf{ in } exp_1, \ldots, x_n \textbf{ in } exp_n \textbf{ with } exp_g\}$$

we refer to the free variable x_0 as the *specifier variable*, the expression exp_0 as the *specifier expression*, the free variables $x_1 \ldots x_n$ as the *constrainer variables*, $exp_1 \ldots exp_n$ as the *constrainer expression*, and exp_g as the *guard*.

The simplest variant of set comprehension binds the specifier variable to a constrainer expression producing a single enumerable element:

Set Plugin: Set Comprehension

$(\!|\{\ ^{\alpha}x \mid {}^{\beta_1}x_1 \textbf{ in } {}^{\gamma_1}\boxed{?}_1\}|\!) \quad \rightarrow$

```
if x = x1 then
  if ¬evaluated(γ1) then
    pos := γ1
  else
    if enumerable(value(γ1)) then
      let s = {m | m ∈ enumerate(value(γ1))} in
        [[pos]] := (undef, undef, setElement(s))
    else
      Error('Free variables may only be bound to enumerable elements')
else
  Error('Constrainer variable must have same name as specifier variable')
```

Notice how we use the *setElement*(s) mapping to get a Set element representation of the set s. This variant would support set comprehension expressions of the form $\{x \mid x \textbf{ in } X\}$ where X is an enumerable element.

A slightly more complex version supports set comprehensions of the form

$$\{x \mid x \textbf{ in } X, y_1 \textbf{ in } Y_1, \ldots, y_n \textbf{ in } Y_n \textbf{ with } \varphi\}$$

where X and Y_i's are enumerable elements and x and y_i's are free variables in φ. This form binds multiple constrainer variables to multiple constrainer expressions, and adds more fine grained control with a guard. The semantic definition of this form involves creating temporary logical variables for each constrainer variable and iterating their values over the values offered by their corresponding constrainer expressions and evaluating the guard for each

combination of these values. A formal semantic definition is provided in Appendix A.5.4. This variant supports set comprehension expressions such as:

$$\{x \mid x \textbf{ in } X \textbf{ with } x > z\}$$
$$\{x \mid x \textbf{ in } \{1,3,5\},\ z \textbf{ in } \{2,4,6\} \textbf{ with } (x+z) \textbf{ in } \{3,4,5,6,7,8,9,10\}\}$$

Finally the most complex variant of the form

$$\{x \textbf{ is } e \mid x_1 \textbf{ in } X_1, \ldots, x_n \textbf{ in } X_n \textbf{ with } \varphi\}$$

in which e is an expression, φ is a guard and x_1 to x_n are free variables in both e and φ, allows the specifier to be defined in terms of a specifier expression. In this form the constrainer variables are themselves expected to be present in the specifier expression, and this expression is re-evaluated for all possible combinations of the contrainer variables. Similar to the previous form, the semantics definition of this form also involves creating logical variables for each constrainer variable, evaluating the guard for each combination of their values, and additionally evaluating the specifier expression for each combination that satisfies the guard. The semantics of this variation is also available in Appendix A.5.4.

The last variation is the most expressive form as it allows the user to create sets using a function on constrainer variable values rather than simply being bound to some subset of a single constrainer expression. Here are two examples of defining sets using this form:

$$\{x \textbf{ is } \{a,b,c\} \mid a \textbf{ in } 1..100,\ b \textbf{ in } 1,2,3,\ c \textbf{ in } aSet\}$$
$$\{x \textbf{ is } y * z \mid y \textbf{ in } \{1,3,5\},\ z \textbf{ in } \{2,4,6\} \textbf{ with } (y+z) \textbf{ in } \{3,4,5,6,7,8,9,10\}\}$$

Operators

The Set plugin extends the vocabulary of the CoreASM engine by providing the following operators: $\subset$, $\cup$, $\cap$, and $\setminus$ (set difference). Here, we present the formal definition of $\subset$ and $\cup$ and refer to Appendix A.5.4 for the definition of the other two operators.

Set Plugin : Operators

$$(\!|\ ^{\alpha}\boxed{?} \subset {}^{\beta}\boxed{?}\ |\!)_{[700]} \quad \rightarrow$$

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if $enumerable(value(\alpha)) \wedge enumerable(value(\beta))$ **then**
    let $lv = enumerate(value(\alpha)), rv = enumerate(value(\beta))$ **in**
      $[\![pos]\!] := (undef, undef, (\forall e \in lv\ e \in rv))$

$(\!|\, {}^{\alpha}\boxed{?} \cap {}^{\beta}\boxed{?} \,|\!)_{[675]}$ → **choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
$pos := \lambda$
ifnone
if SETELEMENT(l) $\wedge$ SETELEMENT(r) **then**
let $l = value(\alpha), r = value(\beta)$ **in**
let $v = \{x \mid x \in enumerate(l) \wedge x \in enumerate(r)\}$ **in**
$[\![pos]\!] := (undef, undef, setElement(v))$

Notice that the evaluation of an operation results in a new Set element rather than modification of an existing Set element.

Aggregation Algorithm

The Set plugin is responsible for the aggregation of update instructions with *setAddAction* and *setRemoveAction* that add or remove elements to and from Set elements. The result of aggregation of set updates on a location will be a regular update assigning a new Set element (representing all the changes) to that location.

For every location with a set partial update, the Set plugin first checks the consistency of update instructions before performing the aggregation. The following requirements informally define the consistency of set update instructions [99]:

- If there is a regular update to a given location l along with partial updates:
  - All regular updates to l may only result in a Set element.
  - There cannot exist two regular updates to l resulting in two different values; this is a typical consistency requirement of regular updates.
  - The Set element S assigned by the regular update(s) on l must satisfy all the add and remove update instructions to l; i.e., $\forall \langle l, v_a, setAddAction \rangle \in updates,\ v_a \in S$ and $\forall \langle l, v_r, setRemoveAction \rangle \in updates,\ v_r \notin S$.
- If there are only partial updates to a given location l:
  - There cannot exist two update instructions adding and removing the same element e to location l.
  - The value of location l in the current state of the simulated machine must be a Set element.

The following rule defines the aggregation algorithm offered by the Set plugin; we have

$$aggregatorRule(setPlugin) \equiv @\mathsf{Aggregate}_{Set}$$

Set Plugin

Aggregate$_{Set}(uMset) \equiv$
  local *resultantUpdate* **in**
  seq
    result := {}
  next
    forall $l \in locsToAggregate$ **do**
      if *regularUpdatesExist* **then**
        if $inconsistentRegularUpdates \vee regularUpdateIsNotSet \vee addRemoveConflictWithRU$ **then**
          HandleInconsistentAggregation$(l, uMset, setPlugin)$
        else
          let *resultantUpdate* = GetRegularUpdate$(l, uMset)$ **in**
            add *resultantUpdate* **to result**
      else
        if $addRemoveConflict \vee setNotInLocation$ **then**
          HandleInconsistentAggregation$(l, uMset, setPlugin)$
        else
          let *resultantUpdate* = BuildResultantUpdate$(l, uMset)$ **in**
            add *resultantUpdate* **to result**

where

$locsToAggregate \equiv \{l \mid \langle l, v, a\rangle \in uMset \wedge a \in \{setAddAction, setRemoveAction\}\}$

$regularUpdatesExist \equiv \exists\langle l, v, updateAction\rangle \in uMset$

$inconsistentRegularUpdates \equiv \exists\langle l, v_1, updateAction\rangle \in uMset,$
$\exists\langle l, v_2, updateAction\rangle \in uMset, v_1 \neq v_2$

$regularUpdateIsNotASet \equiv \exists\langle l, v, updateAction\rangle \in uMset, bkg(v) \neq$ "Set"

$addRemoveConflictWithRU \equiv addConflictWithRU \vee removeConflictWithRU$

$addConflictWithRU \equiv \exists\langle l, v_u, updateAction\rangle \in uMset,$
$\exists\langle l, v_a, setAddAction\rangle \in uMset,\ v_a \notin enumerate(v_u)$

$removeConflictWithRU \equiv \exists\langle l, v_u, updateAction\rangle \in uMset,$
$\exists\langle l, v_r, setRemoveAction\rangle, \in uMset,\ v_r \in enumerate(v_u)$

$addRemoveConflict \equiv \exists\langle l, v, setAddAction\rangle \in uMset, \exists\langle l, v, setRemoveAction\rangle \in uMset$

$setNotInLocation \equiv bkg(getValue(l)) \neq$ "Set"

In the case where at least one regular update exists for a location, after checking the

consistency of partial updates with the regular updates on that location, one of the regular updates will be chosen as the result of the aggregation.

Set Plugin

GetRegularUpdate$(loc, uMset) \equiv$
  choose $u \in uMset$ **with** $uiLoc(u) = loc \wedge uiAction(u) = updateAction$ **do**
    result $:= u$
  forall $u \in uMset$ **with** $uiLoc(u) = loc$ **do**
    $aggStatus(u, setPlugin) := successful$

When there is no regular update for a location, all the partial updates are aggregated into a regular update assigning a new Set element to the location resulting from the addition and removal of elements from the value of the location in the current state.

Set Plugin

BuildResultantUpdate$(l, uMset) \equiv$
  local $newSet$ $[newSet := \{\}]$ **in**
    seq
      forall $e \in enumerate(getValue(l))$ **do**
        if $\nexists \langle l, e, setRemoveAction\rangle \in uMset$ **then**
          add e **to** $newSet$
      forall $\langle l, v, setAddAction\rangle \in uMset$ **do**
        add v **to** $newSet$
    next
      result $:= \langle l, setElement(newSet), updateAction\rangle$
      forall $u \in uMset$ **with** $uiLoc(u) = l$ **do**
        $aggStatus(u, setPlugin) := successful$

Composition Algorithm

The Set plugin provides the semantics of sequential composition of Set partial updates. There are five cases to be considered:

1. If the location is not updated in the second step, all the updates of the first step are carried forward.

2. If the location is not updated in the first step, all the updates of the second step are carried forward.

3. If there is a regular update on the location in the second step (i.e., a Set element is assigned to the location in the second step), all the updates in the first step are discarded and the updates of the second step are carried forward.

4. If there is a regular update on the location in the first step and there are partial updates in the second step, the updates need to be aggregated into one regular update.

5. If there are only partial updates on the location in both the first and the second step, those partial updates in the first step that are overridden by the updates in the second step must be removed.

The Set composition algorithm, capturing the five cases above, is formally defined as follows:

Set Plugin

Compose$_{Set}(uMset_1, uMset_2) \equiv$
  seq
   result := {}
  next
   forall $l \in locsAffected$ **do**
    if $locHasAddRemove(uMset_1) \wedge \neg locUpdated(uMset_2)$ **then**
     forall $ui \in uMset_1$ **with** $uiLoc(ui) = l$ **do**
      add ui **to result**
    else if $\neg locUpdated(uMset_1) \wedge locHasAddRemove(uMset_2)$ **then**
     forall $ui \in uMset_2$ **with** $uiLoc(ui) = l$ **do**
      add ui **to result**
    else if $locHasAddRemove(uMset_2) \wedge locRegularUpdate(uMset_2)$ **then**
     forall $ui \in uMset_2$ **with** $uiLoc(ui) = l$ **do**
      add ui **to result**
    else if $locHasAddRemove(uMset_2) \wedge locRegularUpdate(uMset_1)$ **then**
     add SetAggregateLocation$(l, uMset_1, uMset_2)$ **to result**
    else if $locHasAddRemove(uMset_1) \wedge locHasAddRemove(uMset_2)$ **then**
     forall $ui \in$ EradicateConflictingUpdates$(l, uMset_1, uMset_2)$ **do**
      add ui **to result**

where

$locsAffected \equiv \{l_1 \mid \langle l_1, v, a\rangle \in uMset_1\} \cup \{l_2 \mid \langle l_2, v, a\rangle \in uMset_2\}$

$locHasAddRemove(uMset) \equiv \exists \langle l, v, a\rangle \in uMset, a \in \{setAddAction, setRemoveAction\}$

$locRegularUpdate(uMset) \equiv \exists \langle l, v, a\rangle \in uMset, a = updateAction$

$locUpdated(uMset) \equiv \exists \langle l, v, a\rangle \in uMset$

In case (4), the regular update produced is created by aggregating the partial updates in the second step, assuming that the location currently contains the value of the regular update from the first step. The following rule formally defines the semantics of this aggregation.

Set Plugin

SetAggregateLocation$(loc, uMset_1, uMset_2) \equiv$
  return *resultantUpdate* **in**
    local *newSet* [*newSet* := {}] **in**
      seq
        forall $e \in enumerate(getLocRegularUpdateValue(uMset_1))$
          if $\nexists \langle loc, e, setRemoveAction\rangle \in uMset_2$ **do**
            add e **to** *newSet*
        forall $\langle loc, v, setAddAction\rangle \in uMset_2$ **do**
          add v **to** *newSet*
      next
        $resultantUpdate := \langle loc, setElement(newSet), updateAction\rangle$

where

$getLocRegularUpdateValue(uMset) \equiv$ v *s.t.* $\langle loc, v, a\rangle \in uMset \wedge a = updateAction$

Partial update instructions occurring in a sequence may nullify one another. In case (5), we remove the updates that fall into one of these categories:

- For any location, addition of an element e in the first step followed by the removal of the same element e in the second step, clearly causes no change to the resulting Set element. Update instructions containing both these opposing actions on the same location are removed from the composed update multiset.
- For any location, removal of an element e in the first step is neutralized by the addition of the same element e in the second step. Thus, such removal update instructions should be excluded from the composed update multiset.

The following rule formally defines the composition behavior in case (5):

Set Plugin

EradicateConflictingSetUpdates$(loc, uMset_1, uMset_2) \equiv$
  return *remainingUpdates* **in**
    seq
      $remainingUpdates := \{\!|\}$
    next
      forall $v \in locValues$ **do**
        if $locValAct(uMset_1, v, setAddAction) \wedge locValAct(uMset_2, v, setRemoveAction)$ **then**
          skip
        else if $locValAct(uMset_1, v, setRemoveAction) \wedge locValAct(uMset_2, v, setAddAction)$ **then**
          forall $ui \in \{\!|\langle loc, v, setAddAction\rangle \in uMset_2|\!\}$ **do**
            add ui **to** *remainingUpdates*
        else
          forall $ui \in getAllLocValUpdates$ **do**
            add ui **to** *remainingUpdates*
where
  $locValues \equiv \{v_1 \mid \langle loc, v_1, a_1\rangle \in uMset_1\} \cup \{v_2 \mid \langle loc, v_2, a_2\rangle \in uMset_2\}$
  $locValAct(uMset, v, a) \equiv \exists \langle loc, v, a\rangle \in uMset$
  $getAllLocValUpdates \equiv \{\langle loc, v, a_1\rangle \in uMset_1\} \cup \{\langle loc, v, a_2\rangle \in uMset_2\}$

6.3.3 The Bag Plugin

The Bag plugin extends the CoreASM language with the background of finite *Bags* or multisets. The background of Bag elements (or Multiset elements) is defined by $bagBack \in$ BackgroundElement; we have

$$name(bagBack) = \text{“BAG”}$$
$$newValue(bagBack) = emptyBag$$

We model Bag elements as values of a domain BagElement. The following functions define the interface of Bag elements and provide a mapping between Bag elements and the multisets of elements they represent:

- $bagElement$: Multiset(Element) $\mapsto$ BagElement
  for every multiset of elements, returns a bag element representation of that multiset.
- $bagElement^f$: (Element $\mapsto \mathbb{N}$) $\mapsto$ BagElement

for every mapping of elements to positive integers (multiplicity function), returns a bag element with the given multiplicity function.

- $bagValue$: BagElement $\mapsto$ Multiset(Element)
  for every bag element, returns the multiset of elements that the bag represents.
- $bagMultiplicity$: BagElement $\mapsto$ (Element $\mapsto \mathbb{N}$)
  for every bag element, returns the multiplicity function of the multiset it represents. The value of this function is zero for all the elements that are not in the bag.
- $bagDomain$: BagElement $\mapsto$ Set(Element)
  for every bag element, returns the set of all the elements that are in the bag.

For all $b \in$ BagElement we have:

- $bkg(b) :=$ "Bag".
- $\forall b' \in$ BagElement $\quad equal_{Bag}(b, b') \equiv$
  $bagDomain(b) = bagDomain(b')$
  $\wedge \forall e \in bagDomain(b) \quad bagMultiplicity(b)(e) = bagMultiplicity(b')(e)$
- $enumerable(b)$
  All bag elements are enumerable.
- $enumerate_{Bag}(b) = bagValue(b)$.
- $b \in$ FunctionElement
  All bag elements also behave as functions.
- $class_{fe}(b) = static$
- $\forall e \in$ Element $\quad value_{fe}(b, \langle e \rangle) \equiv numberElement(bagMultiplicity(b)(e))$

To facilitate partial updates of Bag elements, the **add/to**-rule and **remove/from**-rule are supported by the Bag plugin (see Section 6.3.1). We have

$$\forall b \in \text{BagElement}\ isModifiableCollection(b)$$

Since incremental updates on bags do not come with much constraints as for sets (due to multiplicity of elements), instead of using different update actions for adding/removing elements to/from bags, Bag plugin uses a more general action, *bagUpdateAction*, with special

values (elements) that also include the actions of adding, removing, or an ordered combination of adding or removing of elements; the latter is useful in composing incremental updates on bags:

$$computeAddUpdate_{Bag}(loc, value) \equiv$$
$$\quad \{\!|\langle loc, bagUpdateElement(\text{“add”}, value), bagUpdateAction\rangle|\!\}$$
$$computeRemoveUpdate_{Bag}(loc, value) \equiv$$
$$\quad \{\!|\langle loc, bagUpdateElement(\text{“remove”}, value), bagUpdateAction\rangle|\!\}$$

Expression Forms

The interpreter is extended with the following Bag enumeration forms:

Bag Plugin

$(\!\|<<\ >>\|\!)$	$\rightarrow$	$[\![pos]\!] := (undef, undef, emptyBag)$
$(\!\|<<\ {}^{\lambda_1}[?]_1, \ldots, {}^{\lambda_n}[?]_n\ >>\|\!)$	$\rightarrow$	**choose** $i \in [1..n]$ **with** $\neg evaluated(\lambda_i)$ $pos := \lambda_i$ **ifnone** **let** $m = \{\!\|value(\lambda_i) \mid i \in [1..n]\|\!\}$ **in** $[\![pos]\!] := (undef, undef, bagElement(m))$

Various forms of bag comprehension similar in syntax and semantics to those of sets (see Section 6.3.2) is also introduced by the Bag plugin.

Operators

Bag plugin provides the following four operators on Bag elements: $\cap$ (multiset intersection), $\setminus$ (multiset difference), $\cup$ (multiset union), and $+$ (multiset join) as defined below:

Bag Plugin

$(\!\|{}^{\alpha}[?] \cap {}^{\beta}[?]\|\!)_{[675]}$	$\rightarrow$	**choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$ $pos := \lambda$ **ifnone** **let** $l = value(\alpha), r = value(\beta)$ **in** **if** BagElement$(l) \wedge$ BagElement(r) **then** **let** $f = \{x \mapsto y \mid x = (bagDomain(l) \cap bagDomain(r))$ $\wedge\ y = min(bagValue(l)(x), bagValue(r)(x))\}$ **in** $[\![pos]\!] := (undef, undef, bagElement^f(f))$

$(\!|\, {}^{\alpha}[?] \backslash {}^{\beta}[?] \,|\!)_{[650]} \rightarrow$ **choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
$pos := \lambda$
ifnone
let $l = value(\alpha), r = value(\beta)$ **in**
if BAGELEMENT$(l) \wedge$ BAGELEMENT(r) **then**
let $f = \{x \mapsto y \mid x \in bagDomain(l) \cup bagDomain(r)$
$\wedge\, y = max(0, bagValue(l)(x) - bagValue(r)(x))\}$ **in**
$[\![pos]\!] := (undef, undef, bagElement^f(f))$

$(\!|\, {}^{\alpha}[?] \cup {}^{\beta}[?] \,|\!)_{[650]} \rightarrow$ **choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
$pos := \lambda$
ifnone
let $l = value(\alpha), r = value(\beta)$ **in**
if BAGELEMENT$(l) \wedge$ BAGELEMENT(r) **then**
let $f = \{x \mapsto y \mid x \in bagDomain(l) \cup bagDomain(r)$
$\wedge\, y = max(bagValue(l)(x), bagValue(r)(x))\}$ **in**
$[\![pos]\!] := (undef, undef, bagElement^f(f))$

$(\!|\, {}^{\alpha}[?] + {}^{\beta}[?] \,|\!)_{[750]} \rightarrow$ **choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
$pos := \lambda$
ifnone
let $l = value(\alpha), r = value(\beta)$ **in**
if BAGELEMENT$(l) \wedge$ BAGELEMENT(r) **then**
let $f = \{x \mapsto y \mid x \in bagDomain(l) \cup bagDomain(r)$
$\wedge\, y = bagValue(l)(x) + bagValue(r)(x)\}$ **in**
$[\![pos]\!] := (undef, undef, bagElement^f(f))$

6.3.4 The List Plugin

The List plugin extends the CoreASM language providing the background of lists (sequence of elements) with corresponding operators and rule forms. We denote the background of List elements by $listBkg \in$ BACKGROUNDELEMENT; we have

$$name(listBkg) = \text{“LIST”}$$
$$newValue(listBkg) = emptyList$$

List elements are values of the domain LISTELEMENT. The following functions define the interface of list elements and provide a mapping between List elements and the sequence of

elements they represent.

- $listElement$: List(Element) $\mapsto$ ListElement
  returns a list element representing the given sequence of elements.

- $listValue$: ListElement $\mapsto$ List(Element)
  returns the sequence of elements that are represented by the given list element,

- $head_{le}$: ListElement $\mapsto$ Element
  $last_{le}$: ListElement $\mapsto$ Element
  return the first and last elements of the list, or undef_e if the list is empty.

- $tail_{le}$: ListElement $\mapsto$ ListElement
  returns the tail of the list excluding its first element, or an empty list if the list has only one element.

- $cons_{le}$: Element $\times$ ListElement $\mapsto$ ListElement
  $cons_{le}(e, l)$ constructs a new list with e as its head and l as its tail.

- $concat_{le}$: ListElement $\times$ ListElement $\mapsto$ ListElement
  $concat_{le}(l_1, l_2) \equiv cons_{le}(head_{le}(l_1), concat_{le}(tail_{le}(l_1), l_2))$

- $listItem_{le}$: ListElement $\times$ $\mathbb{N} \mapsto$ Element
  $listItem_{le}(l, i) \equiv listValue(l)(i)$

- $take_{le}$: ListElement $\times$ $\mathbb{N} \mapsto$ ListElement
  $take_{le}(list, i)$ returns a list element containing the first i elements of $list$ as a list element. The first element of the list is at index 1.

- $drop_{le}$: ListElement $\times$ $\mathbb{N} \mapsto$ ListElement
  $drop_{le}(list, i)$ returns a list element containing what is left after dropping the first i elements of the list $list$. The first element of the list is at index 1.

For every $l \in$ ListElement, we have

- $bkg(l) =$ "List"

- $\forall l' \in$ ListElement $\quad equal_{List}(l, l') \equiv listValue(l) = listValue(l')$

- $enumerable(l)$
  All list elements are enumerable.
- $enumerate_{List}(l) = listValue(l)$.
- $l \in$ FunctionElement
  All list elements also behave as functions.
- $class_{fe}(l) = static$
- $\forall ne \in$ NumberElement $\; value_{fe}(l, \langle ne \rangle) \equiv$
  $$\begin{cases} listItem_{le}(l, numericValue(ne)), & \text{if } listItem_{le}(l, numericValue(ne)) \neq undef; \\ \mathsf{undef}_e, & \text{otherwise.} \end{cases}$$

Every list element is considered to be a modifiable collection, so we have

$$\forall l \in ListElement \; isModifiableCollection(l)$$

However, List plugin does not offer partial updates on List elements; hence, adding and removing elements to and from List elements cannot be done incrementally. As a result, $computeAddUpdate_{List}$ and $computeRemoveUpdate_{List}$ on lists return an update instruction with a regular update action defined as:

$$computeAddUpdate_{List}(loc, value) \equiv$$
$$\{\!|\langle loc, concat_{le}(getValue(loc), listElement(\langle value \rangle)), updateAction \rangle|\!\}$$
$$computeRemoveUpdate_{List}(loc, value) \equiv$$
$$\begin{cases} \{\!|\langle loc, concat_{le}(left, right), updateAction \rangle|\!\}, & \text{if } |indices(getValue(loc))| > 0; \\ \{\!|\}\!\}, & \text{otherwise.} \end{cases}$$

where

$$indices(le) = \{j \mid j \in [1..|listValue(le)|] \wedge listValue(le)(j) = value\}$$
$$left = take_{le}(getValue(loc), m - 1)$$
$$right = drop_{le}(getValue(loc), m)$$
$$m = min(indices(getValue(loc))$$

Expression Forms

The List plugin extends the interpreter to support List comprehension:

List Plugin

$(\!|\,[\ ]\,|\!) \quad \rightarrow$ **let** $newList = newValue(listBkg)$ **in**
$[\![pos]\!] := (undef, undef, newList)$

$(\!|\,[\ {}^{\lambda_1}\boxed{?}_1, \ldots, {}^{\lambda_n}\boxed{?}_n\ ]\,|\!) \quad \rightarrow$ **choose** $i \in [1..n]$ **with** $\neg evaluated(\lambda_i)$
$pos := \lambda_i$
ifnone
let $l = \langle value(\lambda_1), \ldots, value(\lambda_n)\rangle$ **in**
$[\![pos]\!] := (undef, undef, listElement(l))$

To facilitate locating a specific element in a List element, the List plugin also offers the following expression form that searches a List element for the occurrence of an element and returns an index to the element of interest. If there is no such element in the list, the result will be undef_e. If the element appears more than once in the list, one index will be returned non-deterministically.

List Plugin : Search

$(\!|\,\textbf{indexof}\ {}^{\alpha}\boxed{e}\ \textbf{in}\ {}^{\beta}\boxed{e}\,|\!) \rightarrow$ **choose** $\tau \in \{\alpha, \beta\}$ **with** $\neg evaluated(\tau)$
$pos := \tau$
ifnone
let $e = value(\alpha), v = value(\beta)$ **in**
if $v \in$ ListElement **then**
let $l = listValue(v)$ **in**
choose $i \in [1..|l|]$ **with** $l(i) = e$ **do**
$[\![pos]\!] := (undef, undef, numberElement(i))$
ifnone
$[\![pos]\!] := (undef, undef, \mathsf{undef}_e)$

In addition, the following expression forms, return an index to the first and the last occurrence of an element in a list.

List Plugin : Search

$(\!|$ **first indexof** $^{\alpha}\boxed{e}$ **in** $^{\beta}\boxed{e}$ $|\!) \rightarrow$
  choose $\tau \in \{\alpha, \beta\}$ **with** $\neg evaluated(\tau)$
    $pos := \tau$
  ifnone
    let $e = value(\alpha), v = value(\beta)$ **in**
      if $v \in$ ListElement **then**
        let $l = listValue(v)$ **in**
          let $indices = \{j \mid j \in [1..|l|] \wedge l(j) = e\}$ **in**
            if $|indices| > 0$ **then**
              $[\![pos]\!] := (undef, undef, numberElement(min(indices)))$
            else
              $[\![pos]\!] := (undef, undef, \mathsf{undef}_e)$

$(\!|$ **last indexof** $^{\alpha}\boxed{e}$ **in** $^{\beta}\boxed{e}$ $|\!) \rightarrow$
  // Similar to above; replace $min(indices)$ by $max(indices)$.

Operators

The List plugin provides the following concatenation operator on List elements:

List Plugin : Concatenation

$(\!|\, ^{\alpha}\boxed{?} + ^{\beta}\boxed{?} \,|\!)_{[750]} \rightarrow$ **choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  let $l = value(\alpha), r = value(\beta)$ **in**
    if $l \in$ ListElement $\wedge\, r \in$ ListElement **then**
      $[\![pos]\!] := (undef, undef, concat_{le}(l, r))$

Rule Forms

The List plugin extends the interpreter of the engine to provide the following rule forms facilitating *shifting* of List elements one index to the left or right. In shift left, the first element of the list is dropped into the given location. In shift right, the last element of the list is dropped into the given location.

List Plugin

```
(| shift left α□ into β□ |) →
    choose τ ∈ {α, β} with ¬evaluated(τ)
      pos := τ
    ifnone
      if value(α) ∈ ListElement then
        if loc(β) ≠ undef then
          let updates = {|⟨loc(β), head_le(value(α)), updateAction⟩,
                          ⟨loc(α), tail_le(value(α)), updateAction⟩|}
            [[pos]] := (undef, updates, undef)
        else
          Error('Cannot shift list to a non-location.')

(| shift right α□ into β□ |) →
    choose τ ∈ {α, β} with ¬evaluated(τ)
      pos := τ
    ifnone
      if value(α) ∈ ListElement then
        if loc(β) ≠ undef then
          let le = value(α), l = listValue(le) in
            if |l| ≤ 1 then
              let updates = {|⟨loc(β), last_le(le), updateAction⟩,
                              ⟨loc(α), emptyList, updateAction⟩|}
                [[pos]] := (undef, updates, undef)
            else
              let updates = {|⟨loc(β), last_le(le), updateAction⟩,
                              ⟨loc(α), take_le(le, |l| − 1), updateAction⟩|}
                [[pos]] := (undef, updates, undef)
        else
          Error('Cannot shift list to a non-location.')
```

Functions

The List plugin also extends the vocabulary of the engine to provide the following functions defined on List elements:

- `head: LIST -> ELEMENT`
  $value_{fe}(headFunction, \langle l \rangle) = head_{le}(l)$

- `last: LIST -> ELEMENT`
  $value_{fe}(lastFunction, \langle l \rangle) = last_{le}(l)$

- `tail: LIST -> LIST`
  $value_{fe}(tailFunction, \langle l \rangle) = tail_{le}(l)$

- `cons: ELEMENT * LIST -> LIST`
  $value_{fe}(consFunction, \langle e, l \rangle) = cons_{le}(e, l)$

- `nth: LIST * NUMBER -> ELEMENT`
  $value_{fe}(nthFunction, \langle l, i \rangle) = listItem_{le}(l, numricValue(i))$

- `take: LIST * NUMBER -> LIST`
  $value_{fe}(takeFunction, \langle l, i \rangle) = take_{le}(l, numricValue(i))$

- `drop: LIST * NUMBER -> LIST`
  $value_{fe}(dropFunction, \langle l, i \rangle) = drop_{le}(l, numricValue(i))$

- `reverse: LIST -> LIST`
  $value_{fe}(reverseFunction, \langle l \rangle) =$
  $$\begin{cases} emptyList, & \text{if } |listValue(l)| = 0; \\ reverse(l), & \text{otherwise.} \end{cases}$$
  where
  $reverse(l) \equiv l'$ s.t. $\forall_{i \in [1..|listValue(l)|]}\ listItem_{le}(l', i) = listItem_{le}(l, |listValue(l)| - i + 1)$

- `indexes: LIST -> LIST`
  $value_{fe}(indexesFunction, \langle l \rangle) = listElement(\langle 1, \ldots, |listValue(l)| \rangle)$

- `indices: LIST -> LIST`
  same as `indexes`.

- `setnth: LIST * NUMBER * ELEMENT -> LIST`
  $value_{fe}(setnthFunction, \langle l, n, e \rangle) =$
  $$\begin{cases} l' \text{ s.t. } listItem_{le}(l', numericValue(n)) = e, & \text{if } 1 \leq n \leq |listValue(l)|; \\ \mathsf{undef}_e, & \text{otherwise.} \end{cases}$$

6.3.5 The Queue Plugin

The Queue plugin does not provide any new type domain but it provides two rule forms that operate on Lists elements as queues: **enqueue** and **dequeue**. The former adds an element to end of the list, and the latter removes an element from the head of the list. We present here a formal definition of these two rule forms:

Queue Plugin

$(\!\|$ **enqueue** $^{\alpha}\square$ **into** $^{\beta}\square \|\!)$	$\rightarrow$	$pos := \beta$

$(\!\|$ **enqueue** $^{\alpha}\square$ **into** $^{\beta}l \|\!)$	$\rightarrow$	**if** $value(\beta) \in$ ListElement **then**
		  $pos := \alpha$
		else
		  Error('Cannot enqueue into a non-list.')

$(\!\|$ **enqueue** $^{\alpha}v$ **into** $^{\beta}l \|\!)$	$\rightarrow$	**let** $newList = concat_{le}(value(\beta), listElement(\langle v \rangle))$ **in** $[\![pos]\!] := (undef, \{\!\{\langle l, newList, updateAction\rangle\}\!\}, undef)$
$(\!\|$ **dequeue** $^{\alpha}\square$ **from** $^{\beta}\square \|\!)$	$\rightarrow$	$pos := \beta$

$(\!\|$ **dequeue** $^{\alpha}\square$ **from** $^{\beta}l_2 \|\!)$	$\rightarrow$	**if** $value(\beta) \in$ ListElement **then**
		  if $\|listValue(value(\beta))\| > 0$ **then**
		  $pos := \alpha$
		  else
		  Error('Cannot dequeue from an empty queue.')
		else
		  Error('Cannot dequeue into a non-list.')

$(\!\|$ **dequeue** $^{\alpha}l_1$ **from** $^{\beta}l_2 \|\!)$	$\rightarrow$	**let** $u_1 = \langle l_1, head_{le}(value(\beta)), updateAction\rangle$,
		  $u_2 = \langle l_2, tail_{le}(value(\beta)), updateAction\rangle$ **in**
		  $[\![pos]\!] := (undef, \{\!\{u_1, u_2\}\!\}, undef)$

6.3.6 The Stack Plugin

Similar to the Queue plugin introduced above, the Stack plugin also does not provide any new type domain but it provides two rule forms that operate on Lists as stacks: **push** and **pop**. The former one, pushes an element at the head of a list and the latter one removes

the first element of the list.

Stack Plugin

$(\!\|$ **push** $^{\alpha}e$ **into** $^{\beta}\square$ $\|\!)$	$\rightarrow$	$pos := \beta$

$(\!\|$ **push** $^{\alpha}e$ **into** $^{\beta}l$ $\|\!)$	$\rightarrow$	**if** $value(\beta) \in$ ListElement **then**
		  $pos := \alpha$
		else
		  Error('Cannot push into a non-list.')

$(\!\|$ **push** $^{\alpha}v$ **into** $^{\beta}l$ $\|\!)$	$\rightarrow$	**let** $newList = cons_{le}(v, value(\beta))$ **in** $[\![pos]\!] := (undef, \{\!\|\langle l, newList, updateAction\rangle\|\!\}, undef)$
$(\!\|$ **pop** $^{\alpha}\square$ **from** $^{\beta}l$ $\|\!)$	$\rightarrow$	$pos := \beta$

$(\!\|$ **pop** $^{\alpha}\square$ **from** $^{\beta}l_2$ $\|\!)$	$\rightarrow$	**if** $value(\beta) \in$ ListElement **then**
		  if $\|listValue(value(\beta))\| > 0$ **then**
		  $pos := \alpha$
		  else
		  Error('Cannot pop from an empty stack.')
		else
		  Error('Cannot pop from a non-list.')

$(\!\|$ **pop** $^{\alpha}l_1$ **from** $^{\beta}l_2$ $\|\!)$	$\rightarrow$	**let** $u_1 = \langle l_1, head_{le}(v), updateAction\rangle$,
		  $u_2 = \langle l_2, tail_{le}(v), updateAction\rangle$ **in**
		  $[\![pos]\!] := (undef, \{\!\|u_1, u_2\|\!\}, undef)$

6.3.7 The Map Plugin

The Map plugin extends CoreASM by providing the background of Map elements and the corresponding operators and rule forms defined on them. The background of map elements is denoted by $mapBkg \in$ BackgroundElement; we have

$$name(mapBkg) = \text{"MAP"}$$
$$newValue(mapBkg) = emptyMap$$

Map elements are values of the domain MapElement. The following functions define the interface of map elements and provide a mapping between Map elements to the unary functions or sets of pairs they represent:

- *mapElement* : (ELEMENT $\mapsto$ ELEMENT) $\mapsto$ MAPELEMENT
  returns a map element representing the given mapping of elements to elements.

- *mapElementFromPairs* : SET(LISTELEMENT) $\mapsto$ MAPELEMENT
  if the given set consists of pairs of elements (lists of size two) of the form $[k_i, v_i]$ such that $\forall [k_i, v_i] \not\exists [k_j, v_j] \quad k_i = k_j \wedge v_i \neq v_j$, this function returns a map element representing a mapping of k_is to v_is; otherwise, returns undef_e.

- *mapValue* : MAPELEMENT $\mapsto$ (ELEMENT $\mapsto$ ELEMENT)
  returns the mapping (from elements to elements) represented by the given map element.

- *keyset* : MAPELEMENT $\mapsto$ SET(ELEMENT)
  $\forall m \in$ MAPELEMENT, $keyset(m) \equiv domain(mapValue(m))$

- *valueset* : MAPELEMENT $\mapsto$ SET(ELEMENT)
  $\forall m \in$ MAPELEMENT, $valueset(m) \equiv range(mapValue(m))$

For every $m \in$ MAPELEMENT, we have

- $bkg(m) =$ "Map"

- $\forall m' \in$ MAPELEMENT $\quad equal_{Map}(m, m') \equiv$
  $keyset(m) = keyset(m') \wedge \forall e \in keyset(m) \quad mapValue(m')(e) = mapValue(m)(e)$

- *enumerable*(m)
  All map elements are enumerable.

- $enumerate_{Map}(m) = \{listElement(\langle k, v \rangle) \mid k \in keyset(m) \wedge v = mapValue(m)(k)\}$.

- $m \in$ FUNCTIONELEMENT
  All map elements also behave as functions.

- $class_{fe}(m) = static$

- $\forall e \in$ ELEMENT $\quad value_{fe}(m, \langle e \rangle) \equiv$
  $$\begin{cases} mapValue(m)(e), & \text{if } mapValue(m)(e) \neq undef; \\ \mathsf{undef}_e, & \text{otherwise.} \end{cases}$$

Expression Forms

The Map plugin extends the interpreter of the CoreASM engine with the following map comprehension forms:

Map Plugin

$$\langle\!|\, \{\texttt{->}\} \,|\!\rangle \;\rightarrow\; [\![pos]\!] := (undef, undef, emptyMap)$$

$$\langle\!|\, \{^{\lambda_1}\boxed{?}\texttt{->}\ ^{\lambda_2}\boxed{?},\ \ldots,\ ^{\lambda_{2n-1}}\boxed{?}\texttt{->}\ ^{\lambda_{2n}}\boxed{?}\} \,|\!\rangle \;\rightarrow$$
$$\textbf{choose } i \in [1..2n] \textbf{ with } \neg evaluated(\lambda_i)$$
$$\quad pos := \lambda_i$$
$$\textbf{ifnone}$$
$$\quad \textbf{let } pairs = \{listElement(\langle \lambda_{2i-1}, \lambda_{2i} \rangle) \mid i \in [1..n]\} \textbf{ in}$$
$$\quad\quad [\![pos]\!] := (undef, undef, mapElementFromPairs(pairs))$$

Functions

The vocabulary of the CoreASM engine is also extended with the following two functions mapping Map elements to sets of pairs and vice versa:

- `toMap: ELEMENT -> MAP`

  $$value_{fe}(toMapFunction, \langle e \rangle) = \begin{cases} mapElementFromPairs(\{x \mid x \in enumerate(e)\}), & \text{if } enumerable(e); \\ \mathsf{undef}_e, & \text{otherwise.} \end{cases}$$

- `mapToPairs: MAP -> SET`

  $$value_{fe}(mapToPairsFunction, \langle m \rangle) = \begin{cases} setElement(enumerate(m)), & \text{if } m \in \textsc{MapElement}; \\ \mathsf{undef}_e, & \text{otherwise.} \end{cases}$$

6.4 Auxiliary Plugins

In addition to the plugins addressed so far, CoreASM comes with a number of auxiliary plugins that extend the kernel of CoreASM with concepts, constructs and functionalities that are particularly useful in execution and analysis of specifications. Here, we present those auxiliary plugins that are available as part of the current edition of CoreASM.

6.4.1 The Signature Plugin

The CoreASM language is in principle an untyped language.[4] While a typeless language is desirable for writing initial specifications, defining the types of values and the signatures of functions used in more concrete specifications often add useful semantic information. Such information not only can improve the understandability of the specification and reduce specification errors, but it also plays an essential role in the verification process.

The Signature plugin extends the CoreASM language with syntactic patterns to declare universes, enumerated backgrounds, and function signatures. The corresponding nodes in the parse tree are processed by the Signature plugin when the CoreASM engine is initializing the Abstract Storage (see *Initializing State* in Figure 4.5). During this phase, the engine queries plugins for their contributions to the vocabulary of the state (see definition of InitAbstractStorage in Section 5.5). When the Signature plugin is asked for its vocabulary contribution, it processes the parse tree and provides the engine with a list of universes, backgrounds and functions declared in the specification. Thus, the interpretation of Signature plugin declarations directly modifies the initial state of the simulated machine.

Functions

To declare functions, the Signature plugin extends the CoreASM language with the following syntactic patterns:

Signature Plugin

$$\langle\!|\ \textbf{function}\ x : \texttt{->}\ x_r\ |\!\rangle \quad \rightarrow \quad \mathsf{CreateFunction}(x, \mathit{controlled}, \langle\rangle, x_r)$$
$$\langle\!|\ \textbf{function controlled}\ x : \texttt{->}\ x_r\ |\!\rangle \quad \rightarrow \quad \mathsf{CreateFunction}(x, \mathit{controlled}, \langle\rangle, x_r)$$
$$\langle\!|\ \textbf{function static}\ x : \texttt{->}\ x_r\ |\!\rangle \quad \rightarrow \quad \mathsf{CreateFunction}(x, \mathit{static}, \langle\rangle, x_r)$$

$$\langle\!|\ \textbf{function}\ x : x_{d_1} \texttt{*}\ \ldots \texttt{*}\ x_{d_n} \texttt{->}\ x_r\ |\!\rangle \rightarrow \mathsf{CreateFunction}(x, \mathit{controlled}, \langle x_{d_1}, \ldots, x_{d_n}\rangle, x_r)$$
$$\langle\!|\ \textbf{function controlled}\ x : x_{d_1} \texttt{*}\ \ldots \texttt{*}\ x_{d_n} \texttt{->}\ x_r\ |\!\rangle \rightarrow \mathsf{CreateFunction}(x, \mathit{controlled}, \langle x_{d_1}, \ldots, x_{d_n}\rangle, x_r)$$
$$\langle\!|\ \textbf{function static}\ x : x_{d_1} \texttt{*}\ \ldots \texttt{*}\ x_{d_n} \texttt{->}\ x_r\ |\!\rangle \rightarrow \mathsf{CreateFunction}(x, \mathit{static}, \langle x_{d_1}, \ldots, x_{d_n}\rangle, x_r)$$

The interpretation of function declaration patterns is defined by the CreateFunction rule, which creates a new function with a specified name, class, and signature.

[4] This section is based on Section 5.2 of George Ma's M.Sc. thesis [93] and Section 3.1 of our previously published paper on "Model Checking CoreASM Specifications" [59].

Signature Plugin

CreateFunction$(name, functionClass, domain, range) \equiv$
  let $f = new(\text{FUNCTIONELEMENT})$ **in**
    $class_{fe}(f) := functionClass$
    let $s = new(\text{SIGNATURE})$ **in**
      $sigDomain(s) := domain$
      $sigRange(s) := range$
      $signature(f) := s$
      add $(name, f)$ **to** $pluginFunctions(signaturePlugin)$

One can also specify the initial value(s) of a function in the function declaration by including an initialization expression at the end of the declaration. The initialization expression may be a basic expression, for nullary functions, or a function expression, for n-ary functions. Before the function is created, the expression giving its initial value is evaluated. In the following patterns x_c is either *static* or *controlled*.

Signature Plugin

$(\!|$ **function** x_c x : -> x_r **initially** $^{\alpha}\boxed{e}$ $|\!) \rightarrow$ $evaluate(\alpha)$
$(\!|$ **function** x_c $x : x_{d_1}$ * ... * x_{d_n} -> x_r **initially** $^{\alpha}\boxed{e}$ $|\!) \rightarrow$ $evaluate(\alpha)$

$(\!|$ **function** x_c x : -> x_r **initially** $^{\alpha}v|\!) \rightarrow$ CreateFunctionWithInitValue$(x, x_c, \langle\rangle, x_r, v)$
$(\!|$ **function** x_c $x : x_{d_1}$ * ... * x_{d_n} -> x_r **initially** $^{\alpha}v|\!) \rightarrow$
  CreateFunctionWithInitValue$(x, x_c, \langle x_{d_1}, \ldots, x_{d_n}\rangle, x_r, v)$

CreateFunctionWithInitValue$(name, functionClass, domain, range, initialValue) \equiv$
  let $f = new(\text{FUNCTIONELEMENT})$ **in**
    $class_{fe}(f) := functionClass$
    let $s = new(\text{SIGNATURE})$ **in**
      $sigDomain(s) := domain$
      $sigRange(s) := range$
      $signature(f) := s$
      if $initialValue \neq undef$ **then**
        SetFunctionValue$(f, domain, initialValue)$
      add $(name, f)$ **to** $pluginFunctions(signaturePlugin)$

The SetFunctionValue rule sets the initial value of a function. If the function is not nullary and the specified value is a MAPLEMENT, each key in the map is viewed as an argument list

and the value of the function for those arguments is set to the corresponding map value.

Universes and Enumerations

The Signature plugin also extends the CoreASM language with patterns for declaration of universes:

Signature Plugin

$(\!\| \textbf{universe } x \|\!)$	$\rightarrow$	$\mathsf{CreateUniverse}(x, \{\})$
$(\!\| \textbf{universe } x = \{x_{e_1}, \ldots, x_{e_n}\} \|\!)$	$\rightarrow$	$\mathsf{CreateUniverse}(x, \{x_{e_1}, \ldots, x_{e_n}\})$

The second pattern allows the specification writer to declare a universe along with a set of named initial member elements. Of course, a declared universe can still be extended using standard methods, namely by using the **extend** rule, which imports a new element to a universe, or by setting the value of the corresponding universe membership predicate to *true* for a given element.

The universe declaration patterns are interpreted by the CreateUniverse rule, which creates a new universe with the specified name. If initial members are specified, for each member a static function with the given name is also created.

Signature Plugin

CreateUniverse$(name, members) \equiv$
  let $u = new(\textsc{UniverseElement})$ **in**
    add $(name, u)$ **to** $pluginUniverses(signaturePlugin)$
    forall $elementName \in members$ **do**
      let $e = new(\textsc{Element})$ **in**
        $member_{ue}(u, e) := true$
        let $f = new(\textsc{FunctionElement})$ **in**
          add $(elementName, f)$ **to** $pluginFunctions(signaturePlugin)$
          $class_{fe}(f) := static$
          $\mathsf{SetValue}_{\mathsf{fe}}(f, \langle\rangle, e)$

To declare enumerated backgrounds, the Signature plugin provides the following pattern:

Signature Plugin

$(\!\| \textbf{enum } x = \{x_{e_1}, \ldots, x_{e_n}\} \|\!)$	$\rightarrow$	$\mathsf{CreateEnumeration}(x, \{x_{e_1}, \ldots, x_{e_n}\})$

The CreateEnumeration rule is similar in spirit to CreateUniverse, as enumerable backgrounds are analogous to static universes. The rule is defined as follows:

Signature Plugin

CreateEnumeration$(name, members) \equiv$
  let $b = new(\text{EnumerationBackground})$ **in**
    add $(name, b)$ **to** $pluginBackgrounds(signaturePlugin)$
    forall $elementName \in members$ **do**
      let $e = new(\text{Element})$ **in**
        $bkg(e) := name$
        add e **to** $enumMembers(b)$
        let $f = new(\text{FunctionElement})$ **in**
          add $(elementName, f)$ **to** $pluginFunctions(signaturePlugin)$
          $class_{fe}(f) := static$
          $\mathsf{SetValue}_{\mathsf{fe}}(f, \langle\rangle, e)$

We model background elements that are defined using the Signature plugin with values of the domain EnumerationBackground. The following function, defined on Enumeration Background elements, holds the set of elements each such background represents:

$$enumMembers : \text{EnumerationBackground} \mapsto \text{Set}(\text{Element})$$

For all $eb \in$ EnumerationBackground, we have

- $enumerable(eb)$
  All enumeration background elements are enumerable.
- $enumerate_{EnumerationBackground}(eb) \equiv enumMembers(eb)$

Type Checking on Updates

In order to offer runtime type checking on updates, the Signature plugin extends the control flow of the CoreASM engine by registering for the extension points proceeding the aggregation of updates (see Figure 4.8). We have,

$$\forall em \in \text{EngineMode},\ isPluginRegisteredForTransition(signaturePlugin, Aggregation, em)$$
$$pluginExtensionRule(signaturePlugin) = @\mathsf{CheckUpdateSetForTypes}$$

As a result of this registration, when the control flow of the engine moves from the *Aggregation* control state to either *Step Succeeded* or *Step Failed*, the engine calls the CheckUpdateSetForTypes rule of the Signature plugin. This rule goes through the update set and

for every update checks the arguments and the value of the update against the signature of the function it is updating and reports the inconsistencies. The following rules formally define this process.

Signature Plugin

CheckUpdateSetForTypes ≡
  if *engineProperties*("TypeChecking") = "strict" **then**
    forall ⟨*loc*, *val*, *act*⟩ ∈ *updateSet* **do**
      let $f = stateFunction(state, name_{lc}(loc)), sig_f = signature(f)$ **in**
        if $sig_f \neq undef$ **then**
          CheckArguments($args_{lc}(loc)$, $sigDomain(sig_f)$)
          CheckValue(*val*, $sigRange(sig_f)$)

CheckArguments(*args*, *domain*) ≡
  if |*args*| ≠ |*domain*| **then**
    Error('Number of arguments passed do not match the domain of the function.')
  else
    forall *i* ∈ [1..|*domain*|] **do**
      let *universe* = *stateUniverse*(*state*, *domain*(*i*)) **in**
        if $\neg member_{ue}(universe, args(i))$ **then**
          Error('Argument does not match the domain of the function.')

CheckValue(*v*, *range*) ≡
  let *universe* = *stateUniverse*(*state*, *range*) **in**
    if $\neg member_{ue}(universe, v)$ **then**
      Error('Update value does not match the range of the function.')

6.4.2 The Scheduling Policies Plugin

The Scheduling Policies plugin provides two basic policies for scheduling of agents by the Scheduler. In any CoreASM specification, the particular scheduling policy to be used can be configured using the CoreASM engine's properties (see also Appendix A.4):

- *pluginSchedulingPolicy*(*SchedulingPoliciesPlugin*) ≡
$$\begin{cases} allFirstPolicy, & \text{if } engineProperties(\text{"SchedulingPolicies.Policy"}) = \text{"allfirst"}; \\ oneByOnePolicy, & \text{if } engineProperties(\text{"SchedulingPolicies.Policy"}) = \text{"onebyone"}; \\ undef, & \text{otherwise.} \end{cases}$$

- $newScheduleRule(allFirstPolicy) \equiv @\mathsf{NewSchedule}_{allfirst}$
- $newScheduleRule(oneByOnePolicy) \equiv @\mathsf{NewSchedule}_{onebyone}$

All-First Policy

The *all-first* scheduling policy first tries to schedule all the given agents elements together in one batch. Alternative options will be non-deterministic subsets of the given sets of elements. Applied to the scheduling of agents, this policy first suggests the execution of all the agents together and if that fails, it offers various subsets of agents as alternative options.

Scheduling Policies Plugin

$\mathsf{NewSchedule}_{allfirst}(group, set) \equiv$
  result := $cons(set, \langle s \mid s \in \mathcal{P}(set) \backslash \{set\}\rangle)$

One-by-One Policy

The *one-by-one* scheduling policy provides a schedule that comprises of a series of non-deterministically selected single elements. The policy, however, tries to maintain a "fair" set of schedules over a group by keeping a history of the already scheduled elements and trying to avoid re-scheduling of those elements as long as other non-scheduled elements are still available. Applied to the scheduling of agents in a CoreASM simulation, this policy results in a sequential execution of agents.

Scheduling Policies Plugin

$\mathsf{NewSchedule}_{onebyone}(group, set) \equiv$
  if $group \neq undef$ **then**
    if $scheduleHistory_{obo}(group) = undef \vee set \backslash scheduleHistory_{obo}(group) = \emptyset$ **then**
      choose $e \in set$ **do**
        result := $\langle e \rangle$
        $scheduleHistory_{obo}(group) := \{e\}$
    else
      choose $e \in set$ **with** $e \notin scheduleHistory_{obo}(group)$ **do**
        result := $\langle e \rangle$
        add e **to** $scheduleHistory_{obo}(group)$
  else
    choose $e \in set$ **do**
      result := $\langle e \rangle$

6.4.3 IO Plugin

In an open-system view towards modeling, the system operates in a given environment. The environment affects system runs through actions or events and the system can as well affect the environment by its output. In abstract state machines, the interaction between the system (the machine) and the environment is captured through *monitored* (also called *in*), *shared*, and *out* functions. Monitored functions are controlled only by the environment; they are channels through which the machine observes the environment. In a given state, the values of all monitored functions are determined (and do not change) [25]. Out functions are updated only by the machine and they are read-only for the environment. Shared functions are both controlled and read by the machine and the environment.

The IO Plugin utilizes this machine-environment interaction mechanism of ASM and provides two simple channels of communication between a CoreASM machine and its environment: a **print** rule that outputs values to the environment, and an *input* function to get values from the environment. In both cases, textual representations of values are used.

Functions

To facilitate input from the environment, the IO plugin introduces the following monitored function:

- `input: STRING -> STRING`
  $class_{fe}(inputFunction) = monitored$
  For any given value as its argument, this *input* function queries an input value from the environment (presenting the argument as a prompt or key to the input value). Since this is a monitored function, once its value is set for a certain argument (i.e., message) in a computation step, it will not change before the step is completed.

Rule Forms

To provide an output channel for CoreASM specifications, the IO plugin extends the state of the simulated machine by introducing an `output` function (`output: -> String`) which in any given step holds the output of the previous step. Output values are assigned to `output` by **print** rules. Every **print** rule generates a special update instruction with *printAction* to append the a String element to the value of the `output` function. At the end of each

computation step, these special updates will be aggregated into one single update to `output` function.

IO Plugin

$(\!|\ \textbf{print}\ {}^{\alpha}\boxed{e}\ |\!) \rightarrow \quad pos := \alpha$

$(\!|\ \textbf{print}\ {}^{\alpha}v\ |\!) \rightarrow \quad \textbf{let}\ l = (\text{“output”}, \langle\rangle)\ \textbf{in}$
$[\![pos]\!] := (undef, \{\!|\langle l, stringElement(v), printAction\rangle|\!\}, undef)$

Aggregation of Output Messages

In the aggregation phase of every step, print update instructions need to be aggregated into a single regular update to the `output` function. Since the print values are String elements (see Section 6.2.3), and there is no execution order on the print rules that generated these updates, the aggregation of these values can be achieved by concatenation of the values into a single String element in a non-deterministic order. The IO plugin provides the semantics of such aggregation as follows.

IO Plugin

Aggregate$_{IO}(uMset) \equiv$
  seq
    result := *emptyString*
  next
    if *regularUpdatesExist* **then**
      HandleInconsistentAggregation$(l, uMset, ioPlugin)$
    else
      foreach $u \in$ *printActionUpdates* **do**
        result := *concatString*(**result**, *uiValue*(u))
        aggStatus$(u, ioPlugin)$:= *successful*
where
  regularUpdatesExist $\equiv \exists u \in uMset, uiAction(u) = updateAction \wedge uiLoc(u) = (\text{“output”}, \langle\rangle)$
  printActionUpdates $\equiv \{u \mid u \in uMset \wedge uiAction(u) = printAction \wedge uiLoc(u) = (\text{“output”}, \langle\rangle)\}$

Composition of Output Messages

In order to maintain the order of output values in a sequential composition of print updates, the composition algorithm provided by the IO plugin aggregates the output values of the

first and second step and concatenates them together into a single print update instruction on the `output` function. The the output values of the first step are only considered if the second step does not have a regular update on the output location.

IO Plugin

Compose$_{IO}(uMset_1, uMset_2) \equiv$
  local $outputStr$ [$outputStr := emptyString$] **in**
    seq
      if $\neg regularUpdatesExist(uMset_2)$ **then**
        foreach $u \in printUpdates(uMset_1)$ **do**
          $outputStr := concatString(\mathbf{result}, uiValue(u))$
    seq
      foreach $u \in printUpdates(uMset_2)$ **do**
        $outputStr := concatString(\mathbf{result}, uiValue(u))$
    next
      result $:= \langle(\text{“output”}, \langle\rangle), outputStr, printAction\rangle$
where
  $printUpdates(mset) \equiv \{u \mid u \in mset \wedge uiLoc(u) = (\text{“output”}, \langle\rangle) \wedge uiAction(u) = printAction\}$
  $regularUpdatesExist(mset) \equiv \exists u \in mset, uiAction(u) = updateAction \wedge uiLoc(u) = (\text{“output”}, \langle\rangle)$

6.4.4 The Observer Plugin

It is sometimes desirable to have a machine-readable log of the execution of a CoreASM specification for offline analysis and visualization. One argument for such a feature is that it allows for a clear separation of the execution and the analysis. For example, execution of certain specifications may be time-consuming, but once the execution is done, visualization of the run of the system can be done more quickly and repeatedly, if all the updates of interest are recorded.

The Observer plugin monitors the execution of specifications in CoreASM and produces an XML log of the updates that are produced after every computation step. The plugin can be configured so that only the updates on certain locations of interest are recorded. In order to monitor the updates, the plugin registers itself for the extension point where the control flow of the engine switches to the *Step Succeeded* control state (see Figure 4.6). We have,

$$\forall s \in \text{EngineMode},\ isPluginRegisteredForTransition(observerPlugin, s, stepSucceeded)$$

where *observerPlugin* ∈ PLUGIN is the Observer plugin.

At this point in the engine lifecycle (when the control state changes to *Step Succeeded*), the computation step is successfully completed and the updates are applied to the state. The Observer plugin then simply goes through the last set of updates and records an XML log of those updates that modify the locations of interest.

$$pluginExtensionRule(observerPlugin) = @\mathsf{FireOnModeTransition}_{Observer}$$

Observer Plugin

FireOnModeTransition$_{Observer}$(*sourceMode*, *targetMode*) ≡
  if *targetMode* = *stepSucceeded* **then**
    local *xmlElement* [*xmlElement* := *newStepXMLElement*] **in**
      seq
        foreach *u* **in** *updateSet* **with** *uiLoc*(*u*) ∈ *observerLocationsOfInterest* **then**
          AddXMLChildElement(*xmlElement*, *newUpdateXMLElement*(*u*))
      next
        AppendToLog(*xmlElement*)

6.4.5 Math Plugin

In writing executable specification, one may need to have access to various mathematical constants (such as π) or functions (such as the trigonometric functions) as part of the Number background. The Math plugin addresses this requirement by extending the vocabulary of CoreASM states and providing a number of basic mathematical functions. Most of these functions are equivalent of their Java counterparts defined in the Java library package `java.lang.Math`.

In the following, we present a few of these functions as examples. A complete list of Math plugin functions is provided in Appendix A.5.5.

- `abs(v)` returns the absolute value of v.
- `asin(v)` returns the arc sine of an angle, in the range of $-\pi/2$ through $\pi/2$.
- `floor(v)` returns the largest (closest to positive infinity) value that is less than or equal to the argument and is equal to a mathematical integer.
- `log(v)` returns the natural logarithm (base e) of v.

- `max(v1, v2)` returns the greater of two values.
- `min(v1, v2)` returns the smaller of two values.
- `pow(x, y)` returns the value of the first argument raised to the power of the second argument.
- `powerset(set)` computes the powerset of the given set.
- `sum({v1,...,vn}, @f)` returns the sum of a collection of numbers, after applying function `f` to the values in the collection. If there is one non-number in the collection, it returns *undef*.

An Example

As an example, the output of the execution of Program 6.1 is the following:

```
sum( {1, 2, 100} ) = 103
min(51, 43) = 43
asin(0.5) = 30
powerset({1, 2, 3}) = {{}, {3}, {2}, {3, 2}, {1}, {3, 1}, {2, 1}, {3, 2, 1}}
{2, 3} memberof powerset({1, 2, 3} = true
log(e) = 1
{3, 2, 4} is not a member of powerset({1, 2, 3})
sum( {1, 2, 100}, @a ) = 515
'e' = 2.718281828459045
sin(30) = 0.5
```

6.4.6 The Time Plugin

To introduce the notion of time in CoreASM, the Time plugin extends the vocabulary of the state with a nullary monitored function

```
now: -> NUMBER
```

that provides the current time of the system as a numeric value. Although, such a monitored function seems to be all that is basically needed to have the notion of time in CoreASM, future versions of this plugin could introduce various functions to extract date and time components from time values (e.g., day of the week, hours, or minutes) or to produce specific or relative time values, such as *12/May/2009* or *now - two hours*.

```
CoreASM MathPluginExample

use Standard
use Math

init Init

rule Init = {
    program( self ) := @Main
    a(1) := 5
    a(2) := 10
    a(100) := 500
}

rule Main =
    let e = MathE in {
        print "'e' = " + e
        print "log(e) = " + log(e)
        print "sin(30) = " + round( sin( toRadians(30) ) * 10 ) / 10
        print "asin(0.5) = " + round( toDegrees( asin(0.5) ) )
        print "min(51, 43) = " + min(51, 43)
        print "sum( 1, 2, 100 ) = " + sum({1, 2, 100})
        print "sum( 1, 2, 100, @a ) = " + sum({1, 2, 100}, @a)
        print "powerset(1, 2, 3) = " + powerset({1, 2, 3})
        print "2, 3 memberof powerset(1, 2, 3 = "
                    + ({2, 3} memberof powerset({1,2,3}))
        choose x in powerset({1, 2, 3, 4}) do
            if x memberof powerset({1, 2, 3}) then
                print x + " is a member of powerset(1, 2, 3)"
            else
                print x + " is not a member of powerset(1, 2, 3)"
        ifnone
            print powerset({1, 2, 3})
    }
```

Program 6.1: A CoreASM Example Using Math Plugin

6.5 The JASMine Plugin

In this chapter we have introduced various CoreASM plugins implementing most common mathematical objects and structures, such as *numbers*, *sets*, *lists*, and *maps*.[5] While these backgrounds are usually sufficient for modeling most algorithms and systems, complex specifications may need more advanced features, not necessarily data-oriented. For example, an executable specification for a new peer-to-peer protocol may need access to *network sockets* and *files*; a specification that is used as an executable stub for a software module that still has to be implemented or for a missing piece of hardware may need to put up an on-screen *window* showing its current state; a complex numerical algorithm which is already specified by some standard may be moved out of a specification and a *concrete implementation* written in a standard programming language may be used in its place.

There is thus a clear need to allow *interaction* between CoreASM specifications and concrete code, including operating systems functions, external libraries, and custom code. Among the various tools for running ASM models [48], AsmL (ASM Language) [101], XASM (eXtensible ASM) [4], and AsmGofer [113] provide some support for interaction with external programming languages. AsmL, built on the Microsoft .NET framework [100], incorporates numerous object-oriented features and constructs of Microsoft .NET and supports interaction with external .NET classes. The XASM language allows external C-functions to be used in XASM specifications. However, the arguments and return values of C-functions can only be of a specific C-type that represents elements of the super-universe in XASM. Newer versions of XASM support interaction with Java classes but the support is only limited to invoking Java object constructors. AsmGofer [113], an ASM interpreter embedded in the functional programming language "Gofer", supports the use of functional programming in the definition of types and functions.

In this section, we present JASMine, a CoreASM plugin that offers a solution for the interaction of CoreASM specifications and concrete code by integrating Java with CoreASM.

[5] This section is based on a joint work with Dr. Vincenzo Gervasi and is currently under publication in [68].

6.5.1 Requirements and Limitations

The Java Class Library provides an extremely rich (and continuously growing) set of APIs and efficient implementations for almost any computing task. Moreover, Java offers platform-independence, support on a wide variety of architectures, and many modern language features that make it an attractive target for the integration of ASM specifications with concrete code.

However, there is a risk that by intertwining the "ASM world" of elements, functions and predicates and the "object world" of an object-oriented language, the very nature of the ASM paradigm may be changed in fundamental ways. This is, for example, what happened in AsmL [101], where rules and methods, elements and objects, sets and the Set object of the .NET framework become confused.

In contrast to AsmL, we do not want interaction with Java to *pollute* the CoreASM word. In particular,

- we want to maintain typelessness of the language: it must be possible to treat Java objects as regular ASM values, and to pass untyped ASM elements as arguments to Java methods (with type checking performed at run time only);
- we want to maintain the parallel model of execution of ASMs: the notion of *step* must be preserved, as well as the assumption that the ASM state and environment is observed in a stable snapshot, and updates are applied in parallel and only when no conflicts arise;
- we want to avoid the introduction of extraneous fundamental concepts: the notions of *state*, *update* and *step* should suffice to describe the computation.

The fundamental choice of preserving the ASM computation model sets strong constraints on how JASMine works, which will be described later in more detail.

Four basic capabilities are needed for a minimal reasonable level of interaction, namely: 1) instantiating new objects, invoking their constructors, and storing a reference to the new object in the ASM state; 2) accessing (reading and writing) public fields of objects, including static fields of classes; 3) invoking public methods of objects and static methods of classes, passing the needed arguments, and storing the result in the ASM state; 4) converting between certain ASM types and the corresponding Java types and back, as needed to support expression evaluation and updates. The mechanisms we propose to provide these capabilities

constitute a *conservative extension* of CoreASM, in the sense that the semantics of the non-JASMine parts of a specification are not altered by the extension[6].

Notice that the integration that JASMine provides between ASMs and Java is far less complete than the one existing between, for example, AsmL and .NET: in particular, it is not currently possible to define new Java classes or interfaces through ASM specifications, nor is it possible to use Java inheritance in CoreASM specifications. Interfaces and abstract classes cannot be accessed at all.

We do not see these limitations as particularly relevant in practice. In fact, the design goal of JASMine is to allow *interaction* between ASMs and Java, rather than full *integration*, and we believe the JASMine plugin serves well in this capacity.

6.5.2 Language Extensions

The following subsections describe in turn the constructs implementing the four capabilities mentioned above.

Creation of Java Objects

Java objects in JASMine are seen as part of the *environment*, not of the *state*. This is a fundamental design choice, which differs from what others have done (e.g., AsmL), and helps in cleanly separating the structures-based state of ASM, which only changes between steps and through non-conflicting updates, from the independently evolving state of Java, which can change at any time and also due to external events (e.g., a timer or GUI actions).

JASMine introduces a new background (hence, a new kind of element in the ASM state) called `JObject` which holds a *reference* to the real Java object. Only this immutable reference enters the ASM state as a value, and only through a special update command, hence the basic ASM computation cycle is preserved. As a consequence, creation of a new object is not considered an expression (as is the `new` operator in Java) but rather a rule, since it results in an update. We have

$$jObjectBack \in \textsc{BackgroundElement}$$
$$name(jObjectBack) = \text{“JOBJECT”}$$
$$newValue(jObjectBack) = newJObject()$$

[6] In other terms, a specification which does not interact with Java, and thus does not use the JASMine constructs, has the same semantics whether it includes the JASMine plugin or not.

where *newJObject*() returns a new `JObject` element pointing to a new Java object.

In formal terms, using the notation described above, creation of a new Java object is accomplished as follows:

CreationRules

(| **import native** $^{\alpha}\square$ **into** $^{\beta}\boxed{l}$ |) $\rightarrow$ $pos := \beta$

(| **import native** $^{\alpha}x$ **into** $^{\beta}l$ |) $\rightarrow$
```
if isJavaClassName(x) then
  if hasEmptyConstructor(x) then
    EvaluateImport(l, x, ⟨⟩)
  else
    Error('Constructor not found.')
else
  Error('Java class not found.')
```

(| **import native** $^{\alpha}x(^{\lambda_1}\boxed{e}_1, \ldots, ^{\lambda_n}\boxed{e}_n)$ **into** $^{\beta}\boxed{l}$ |) $\rightarrow$ $pos := \beta$

(| **import native** $^{\alpha}x(^{\lambda_1}\boxed{e}_1, \ldots, ^{\lambda_n}\boxed{e}_n)$ **into** $^{\beta}l$ |) $\rightarrow$
```
if isJavaClassName(x) then
  choose i ∈ [1..n] with ¬evaluated(λ_i)
    pos := λ_i
  ifnone
    if hasConstructor(x, ⟨jValue(value(λ_1)), ..., jValue(value(λ_n))⟩)
      EvaluateImport(l, x, ⟨λ_1, ..., λ_n⟩)
    else
      Error('Constructor not found.')
else
  Error('Java class not found.')
```

Here, we use the *jValue* function to abstract from the task of potentially converting CoreASM elements to Java objects (see Page 166). The actual evaluation of the **import native** statement is defined by the following macro, which takes as parameters a location where to store the reference to the new Java object (as a JObject value), an identifier representing the name of the class, and a sequence of positions of values, which will be the actual parameters for the constructor call:

EvaluateImport

EvaluateImport$(l, x, \langle \lambda_i, \ldots, \lambda_n \rangle) \equiv$
  let $u =$ DefUpd(CREATE, $(l, x, \langle jValue(value(\lambda_1)), \ldots, jValue(value(\lambda_n)) \rangle))$ **in**
    let $jtl =$ ("jasmChannel", $\langle\rangle$) **in**
      $[\![pos]\!] := (undef, \{\!|\langle jtl, u, jasmAction \rangle|\!\}, undef)$

Notice in the specification above how the execution of the rule does not really instantiate the new object (whose constructor could have side effects, and thus alter the Java state), but instead accumulates a special update instruction (a *deferred update*) akin to the update instructions used for aggregation and partial updates [99]. Actual instantiation will be performed at update application time, as will be shown later on. The designated location ("jasmChannel", $\langle\rangle$) accumulates all the JASMine-related update instructions that are performed during a step, whereas the DefUpd macro produces an encoding of its parameters, suitable for later execution of the relevant update.

While the subject will be discussed more fully in the following, it is worthwhile to remark here that this strategy ensures that any action that can perturb the environment (e.g., instantiation of a new Java object) will only be taken if the step turns out to be effective, i.e. if no conflicting updates are generated in that step.

Access to Fields of Java Objects

Reading a field in a Java object does not have side effects and thus can be treated as a pure expression as far as the ASM computation cycle is concerned[7]. In particular, the value in the field can be computed immediately at expression evaluation time. In contrast, writing into a field has observable side effects, and thus cannot be performed *during* a step, but only *between* steps; the corresponding value is then stored in the field at update application time through another deferred update. The following rules detail the semantics used for field access in JASMine.

[7] In a multi-threaded context, field values can change at any moment, even without any write action by the ASM specification. To guarantee the stability of the environment, values read from Java fields are cached by JASMine when first read, and the same value is used if the same field read expression on the same Java object is evaluated multiple times in the same step.

FieldReadExpression

$(\!|\ ^{\alpha}\boxed{e}\ \text{->}^{\beta} x\ |\!) \quad \rightarrow \quad pos := \alpha$

$(\!|\ ^{\alpha} v \text{->}^{\beta} x\ |\!) \quad \rightarrow$

if $isJObject(v)$
  if $hasField(jObj(v), x)$
    if $ImplicitConversionMode$ **then**
      $[\![pos]\!] := (undef, undef, asmValue(\mathsf{GetField}(jObj(v), x)))$
    else
      $[\![pos]\!] := (undef, undef, newJObject(\mathsf{GetField}(jObj(v), x)))$
  else
    Error('No such field.')
else
  Error('Not a Java object.')

As can be observed, field access expressions are evaluated by first evaluating the reference to the JObject, and then (after checking that the given value is actually a JObject and that the corresponding class has an accessible field with the given name) the value in the field of the Java object is retrieved, possibly converted to its ASM counterpart based on the configuration of the plugin (see Section 6.5.2), and finally used as the value of the whole expression. Access to static class fields are handled similarly, and we skip here the corresponding rules for brevity.[8] Assignments are treated through deferred updates:

[8]Reading a static field of a class that has a static block and is not initialized can potentially have side effects. Currently, we do not handle this special case and treat static fields and object fields the same with regard to read access.

FieldWriteRule

```
(| store α[e] into β[e] ->γ x |) →
    choose λ ∈ {α, β} with ¬evaluated(λ)
      pos := λ
    ifnone
      if isJObject(value(β)) then
        if hasField(jObj(value(β)), x) then
          let u = DefUpd(STORE, (value(β), x, jValue(value(α)))) in
            let jtl = ("jasmChannel", ⟨⟩) in
              [[pos]] := (undef, {|⟨jtl, u, jasmAction⟩|}, undef)
        else
          Error('No such field.')
      else
        Error('Not a Java object.')
```

Notice how write access to fields is treated as a partial update to the internal structure of the JObject element. Before the engine applies the updates to the state, the JASMine plugin as the corresponding aggregator will have to check that no conflicting assignments to the same field of a given JObject element are performed, and moreover that the JObject as a whole is not updated to a different value in the same step[9]. Once more, write access to static fields of classes is very similar and we do not detail it here.

Invoking Methods of Java Objects

As remarked above, invocation of methods in Java objects can have side effects which can change both the internal state of the object and of other objects as well (i.e., by calling other methods or accessing public fields). For this reason, method invocation is handled through a deferred update, as described below. Two forms of method invocation exists: one for *void* methods, which have no return value, and one for methods returning a value. The simplest version for void methods invocation is specified as follows:

[9]The same situation is found in other cases, e.g. when both $a := \{1, 2\}$ and **add** 3 **to** a appear in the same step.

VoidMethodInvocationRule

$(\!|$ **invoke** $^{\alpha}\square$ ->$^{\beta}x(^{\lambda_1}\square_1, \ldots, ^{\lambda_n}\square_n)|\!) \rightarrow$
  choose $\lambda \in \{\alpha, \lambda_1, \ldots, \lambda_n\}$ **with** $\neg evaluated(\lambda)$
    $pos := \lambda$
  ifnone
    if $isJObject(value(\alpha))$
      if $hasMethod(jObj(value(\alpha)), x, \langle jValue(value(\lambda_1)), \ldots, jValue(value(\lambda_n))\rangle)$
        let $u =$ DefUpd(INVOKE,
                $(undef, value(\alpha), x, \langle jValue(value(\lambda_1)), \ldots, jValue(value(\lambda_n))\rangle))$ **in**
          let $jtl =$ ("jasmChannel", $\langle\rangle$) **in**
            $[\![pos]\!] := (undef, \{\!|\langle jtl, u, jasmAction\rangle|\!\}, undef)$
      else
        Error('No such method.')
    else
      Error('Not a Java object.')

The version for non-*void* methods is only slightly more complex. We provide a special update instruction (in the vein of **add ... to ...**) so that the actual method call is only performed if the update set is guaranteed to be consistent (see section 6.5.2 for detailed conditions).

This solution may be inconvenient at times. For example, it is not possible to assign directly the result of a method invocation to a field of the same or of a different object, as two separate **invoke** and **store** instructions are needed, and in two different steps. In other words, the effect of any rule altering the state of the "Java world" is only observable in the *next* step of the machine, which of course discourages programming in a sequential style: instead, any needed sequentiality will have to be made explicit, e.g. by using an FSM representation of the ASM. Also, field updates and method invocations performed in the same step will be performed—in due time—in an unspecified order, since update instructions in CoreASM constitute an unordered multiset. This behavior, too, may surprise the unaware Java programmer at his first approach with ASMs, as will be discussed in Sections 6.5.4 and 6.5.5.

Nevertheless, we believe that the soundness of the semantics that is given by the deferred updates approach is worth the inconvenience, and can actually help even novice specifiers in drawing a clear line between what needs to be specified and the actual behavior (possibly,

over-specified) of the implementation.

Formally, invocation of non-*void* methods is specified as follows:

NonVoidMethodInvocationRule

```
(| invoke α[e] ->β x(λ1[e]1, ..., λn[e]n) result into γ[l] |) →
    choose λ ∈ {α, γ, λ1, ..., λn} with ¬evaluated(λ)
      pos := λ
    ifnone
      if isJObject(value(α))
        if hasMethod(jObj(value(α)), x, ⟨jValue(value(λ1)), ..., jValue(value(λn))⟩)
          if loc(γ) ≠ undef
            let u = DefUpd(INVOKE, (loc(γ), value(α), x,
                                    ⟨jValue(value(λ1)), ..., jValue(value(λn))⟩)) in
              let jtl = ("jasmChannel", ⟨⟩) in
                [[pos]] := (undef, {|⟨jtl, u, jasmAction⟩|}, undef)
          else
            Error('Cannot update a non-location.')
        else
          Error('No such method.')
      else
        Error('Not a Java object.')
```

As for the previous constructs, we do not detail here how static methods on classes are invoked, as the mechanism is totally analogous.

In practice, if an exception is returned, two updates are produced: one storing the value of the exception (as an ASM JObject) in a designated location, and another one storing a different value to the same location. As a consequence, Java exceptions are mapped in ASMs to conflicting updates, which can be caught via the Turbo ASM **try/catch** rule [25].

Type Conversion

JASMine operates in two type conversion modes: *implicit conversion* and *explicit conversion*. In the implicit mode, which is the default mode, JASMine automatically converts types between CoreASM and Java when needed. This reduces the hassle of type conversion and helps in writing more concise CoreASM specifications. Automatic type conversion, however, has its drawbacks in certain applications: it converts values even when such a conversion is not needed; e.g., when returned values of Java methods are to be passed as arguments in

Java type	**CoreASM background**
bool, Boolean	Boolean
byte, short, int, long, float, double, Byte, Short, Integer, Long, Float, Double	Number
char, Character	currently not supported
String	String
Set interface	Set
List interface	Sequence
Map interface	Function (dynamic)
arrays	currently not supported
any other object	JObject

Table 6.1: Type Conversions Between CoreASM and Java.

future calls to other Java methods. In the explicit mode, the user is responsible for explicitly converting values between Java and CoreASM using the provided CoreASM functions described further below.

JASMine constructs apply type conversion when needed, through the functions *javaValue* and *asmValue* that convert CoreASM values to Java objects and vice versa. These two functions are defined by cases as summarized in Table 6.1. In most of the rules presented in this paper, the *jValue* function abstracts the details of type conversion based on the conversion mode.

The JObject background offers the following two functions, which perform the same conversion on arbitrary values:

- `toJava: Element -> JObject`
  $value_{fe}(toJavaFunction, \langle v \rangle) = javaValue(v)$
- `fromJava: JObject -> Element`
  $$value_{fe}(fromJavaFunction, \langle v \rangle) = \begin{cases} asmValue(jObj(v)), & \text{if } isJObject(v); \\ \mathsf{undef}_e, & \text{otherwise.} \end{cases}$$

Aggregation of Deferred Updates

As we have seen, any modification to the "Java world" is performed through special update instructions, called deferred updates (but not to be confused with ASM updates), to ensure a stable state and a stable environment in course of a single ASM computation step. Three types of deferred updates are used by JASMine: instantiation (CREATE), field writing

(STORE) and method invocation (INVOKE).

Each type of deferred update carries the information necessary for its execution; in particular, CREATE carries information on the Java class to create and on the location of the new ASM element to create; STORE carries information about the JObject whose field is to be modified, about the name of the field to modify, and about the new value to be written in the field; INVOKE carries information about the JObject on which the method has to be invoked, about the name of the method, and about the (possibly empty) list of arguments to pass to the method.

The following compatibility conditions must be met for a set of updates to be considered consistent:

1. No other update is permitted on the ASM location used in a CREATE. Notice that this includes JASMine deferred updates (i.e., it is not possible to import twice to the same location) as well as regular updates (i.e., it is not possible to assign a different value through the assignment operator `:=` or other update rules to a location used in a CREATE).

2. If multiple STOREs are performed on the same field of the same object, they must all assign the same value.

3. Any location used to store the result of an INVOKE cannot appear in any other update.

Notice that this latter condition is sufficient, but not necessary to guarantee consistency. In fact, we disallow even multiple updates that would write the same value (which are normally permitted under standard ASM semantics). The reason for this more restrictive choice is that in general it is impossible to know which value will be returned by a method call without actually calling the method, and we want the method to be called only if a consistent set of updates is generated. Hence, we require a stronger guarantee than what is strictly needed.

If the set of update instructions is consistent, the prescribed operations are performed *in unspecified order*. Notice that the first condition above ensures that newly-created JObjects are not used in the same step, so there is no need to specify a special ordering with CREATE update instructions performed before STORE and INVOKE ones.

A common troublesome case is when multiple method invocations are performed: if the particular sequence is order-sensitive, ordering will have to be specified explicitly by

using a finite state automaton. In most cases, though, the specific order will be immaterial (e.g., `Point.setX()` and `Point.setY()`), and in these cases multiple invocations can well be specified in the same step. We regard this as a desirable feature for a specification: in fact, the implementer will know that fields can be written and that methods can be invoked in any order as long as they are specified to happen in a single ASM step, whereas the ordering between different steps is significant, and should be respected in the implementation.

6.5.3 Implementing JASMine

In its capacity as a bridging technology, JASMine has to interact closely with both the CoreASM engine and the Java virtual machine. We will discuss these interactions in the following.

Interacting with the CoreASM Engine

The CoreASM extensibility architecture dictates that plugins extending the basic CoreASM language have to implement one or more interfaces, depending on which elements of the language (both syntax and semantics) and of the computation cycle are contributed. In particular, JASMine provides the following extensions:

- It implements the *parser plugin* interface to extend the parser with new syntax for native import, field read/write, and method invocation. The syntax rules contributed to the language correspond to the syntactical patterns shown in Section 6.5.1.
- It implements the *interpreter plugin* interface and contributes the semantics for the new syntactical patterns. The semantics contributed correspond to the ASM rules shown in Section 6.5.1.
- It implements the *vocabulary extender* interface to extend the CoreASM state with the JObject background and the monitored *jasmChannel* function. In particular, the two casting functions `toJava` and `fromJava` are introduced as part of the JObject background. Moreover, element equality, ordering and conversion to a String value are forwarded to the Java object represented by any given JObject value.
- It implements the *aggregator* interface to provide aggregation rules which encode all the JASMine update instructions computed in one step into one single update to the *jasmChannel* location.

- To actually communicate with the Java virtual machine, the value of *jasmChannel* must be read after every successful step and the actions encoded therein must be parsed and applied to the corresponding Java objects. To perform this, the JASMine plugin extends the lifecycle of the CoreASM engine and reads the value of *jasmChannel* whenever the control state of the engine is switched to *Step Successful*, i.e. whenever a step is completed with a consistent set of updates; it then executes all the CREATE, STORE and INVOKE operations stored in *jasmChannel*.

Interacting with the JVM

Interaction between JASMine and the Java Virtual Machine is limited to a few, well-defined operations, and is mostly mediated by the Java Reflection API [120].

The application of updates encoded in *jasmChannel* entails the following steps.

1. For CREATE updates, the classical `Class.forName()` method is invoked, passing a string representation of the imported class name. Once a `Class` object for the desired class is obtained, if the nullary version of `import native` was used (i.e., with no arguments passed to the constructor of the object), the `Class.newInstance()` method is invoked to obtain the instance. Otherwise, `Class.getConstructor()` is called to retrieve the corresponding constructor, then the constructor's `newInstance()` method is called, with the given arguments, to obtain the instance. A new JObject element encapsulating the new instance is then created and assigned to the ASM location provided in the CREATE record.

2. For STORE updates, the class of the referenced object is obtained by calling `getClass()` on the reference held by the JObject; the `Field` object is then retrieved through `Class.getField()`, and finally `Field.set()` (or one of its primitive type variants) is called to assign the value from the STORE record.

3. For INVOKE updates, the class of the referenced object is obtained as above, then the matching `Method` object is retrieved through `Class.getMethod()` (notice that in this way only public methods can be retrieved), and finally `Method.invoke()` is called, with the appropriate parameters from the INVOKE record. If the method was non-void, the resulting value is then stored in the ASM location provided in the INVOKE record.

It is worthwhile to remark that fields and methods name resolution is entirely delegated to the Reflection API, and thus follows the normal resolution algorithm in Java (see [72, sections 8.2 & 8.4]).

Evaluation of field read access is performed immediately upon encountering the corresponding expression, by first obtaining the `Field` object as for STORE updates, then invoking `Field.get()` (or one of its primitive types variants) to retrieve the field value, which is then returned as the expression's value. These operations constitute the GetField macro used in the semantics (Section 6.5.2).

The various functions used in Section 6.5.2 (*isJavaClassName*, *hasEmptyConstructor*, *hasConstructor*, *hasField*, *hasMethod*) are directly mapped to the corresponding Reflection API methods. All these predicates are implemented by trying to access the given class, constructor, field or method and possibly catching the various exceptions (`ClassNotFoundException`, `NoSuchMethodException`, `NoSuchFieldException`) thrown by the Reflection API methods. The *jObj* function returns a reference to the Java object encapsulated by a JObject.

Finally the conversion functions *javaValue* and *asmValue* are implemented by cases, as summarized in Table 6.1. In particular, when converting from CoreASM elements to Java values (*javaValue* function), Booleans and numbers are simply converted to the corresponding primitive types in Java; numbers are generally converted to double, then downcast as needed to fit smaller types. CoreASM's strings are wrappers around Java strings, so the conversion is trivial. More complex mathematical structures (e.g., set or sequences) are generally implemented in CoreASM as wrappers to the various Java Collections API objects, so in this case also conversion amounts to unwrapping the underlying object. Any other CoreASM value is upcast to `Object` and passed as-is, thus realizing an opaque container for the ASM value from the point of view of Java code.

Conversion from Java values to CoreASM elements (*asmValue* function) is similar, except that any unrecognized Java object is wrapped in an opaque JObject element from the point of view of ASM code. This allows access to fields and invocation of methods of objects returned from other Java methods, as in

```
invoke calendar->getCurrentDate() result into today
```

followed, in a subsequent step, by

```
wday := today->weekDay
invoke today->add(7) result into nextWeek
```

6.5.4 A Simple Example

In this section, we present a simple example of an ASM using JASMine constructs. Our example, presented in Program 6.2, executes in three steps (distinguished by the `mode` function ranging from 1 to 3) and demonstrates the employment of the sorting capabilities of the standard Java library.

In the first step, we instantiate a `SortedSet` Java object based on a CoreASM list element. Here, JASMine automatically converts the CoreASM list (and all its elements) into their equivalent Java objects. In the second step, three tasks are done in parallel: the resulting `SortedSet` Java object is printed out, its size is retrieved and stored in a CoreASM location (by invoking its `size()` method), and a new value (15) is added to the list. In the last step, the size of the list and its new value (after adding 15) is printed out. Here is the output of execution:

```
The list is [4, 8, 10, 32]
Size of list is 4
After adding 15, the list is [4, 8, 10, 15, 32]
```

Notice that the values of the list are automatically sorted in the `SortedSet` Java object and the order is maintained even after the addition of `15`. It is also interesting to note that since the addition of `15` is done in parallel with retrieving the size of the list, different runs of the specification may result in either of the values 4 or 5 for the size of the list in the output, depending on in which order these two method calls (`size()` and `add(15)`) are performed by JASMine.

6.5.5 Final Remarks

As we mentioned earlier, in defining the semantics of JASMine we have chosen to be faithful to the theoretical ASM model. This choice has important pragmatic implications that we discuss here.

In particular, JASMine presents a *stable view of the Java environment* to ASMs. This is required by ASM semantics, but may be inconvenient in practice, as any action performed on a Java object (e.g., storing a value in a field or invoking a method) will produce observable effects only in the *next* step of the machine: thus, many programming patterns typical of sequential programming cannot be applied. This is also true in the case of Turbo ASM rules: hence, the n-th step in a **seq** or **iterate** rule will *not* observe the effects on the environment

```
CoreASM JASMineExample

use Standard
use Jasmine

function mode:  -> NUMBER initially 1

init InitRule

rule InitRule = {
    case mode of
        1:  import native java.util.TreeSet([8, 10, 4, 32]) into list

        2:  {
            print "The list is " + list
            invoke list->size() result into s
            invoke list->add(15)
        }

        3:  {
            print "Size of list is " + s
            print "After adding 15, the list is " + list
        }
    endcase
    mode:= mode + 1
}
```

Program 6.2: An Example to Illustrate Application of JASMine in CoreASM

of the previous $n-1$ steps, as the corresponding updates are being deferred as described in Section 6.5.2. This is due to the impossibility of rolling back the Java environment to a previous state, which prevents speculative execution of the inner steps of a Turbo ASM step. For example, a **while** cycle like

```
import native java.io.File into file
...
while (lastModified <= lastActed)
    invoke file->lastModified() result into lastModified
...
```

which could be used to wait for a modification to a file, will not work as expected: in fact, invocations to `lastModified()` will be deferred until the end of the step, most probably defeating the programmer's intention.

In terms of style, one could argue that such behavior should be either encapsulated inside a single Java method `waitModification()` (to be invoked through JASMine), or lifted up to the top level of the ASM specification.

Part III

Applications and Conclusions

Chapter 7

Implementing CoreASM

As we addressed in Section 1.4, one of the requirements of the CoreASM modeling environment is that it should be implemented as an open framework, under an open source license, and using a platform-independent programming language, so that it can be later improved or modified as needed by its community of users. Realizing this requirement, we decided to implement CoreASM using the Java programming language, one of the most popular platform-independent[1] programming languages available.

In order to make CoreASM and its source code freely available for both the academic environment and the industry, we had to carefully choose an open source license that provides users and developers the freedom they need to use and modify CoreASM, without the restrictions that come with many open source licenses. After considering various open source licenses such as GNU Lesser General Public License (LGPL) [62], Apache Software License [63], and BSD licenses [109] and looking at similar open source projects, we have decided to make CoreASM source code available under the Academic Free License (AFL) version 3.0[2]. AFL 3.0 is an open source license with no reciprocal obligation to disclose source code; i.e., derivative works can be licensed under other licenses, and the source code of those derivative works need not be disclosed. Such a license provides a good compromise between the availability of the original source code in a free form and the existence of potentially proprietary editions and extensions in the industry.

[1] According to Java's download page on http://java.sun.com, its standard edition is available on a wide variety of hardware and software platforms: Linux, Linux Intel Itanium, Linux x64, Solaris SPARC, Solaris x64, Solaris x86, Windows, Windows Intel Itanium, and Windows x64.

[2] http://www.opensource.org/licenses/afl-3.0.php

Currently, the CoreASM project is publicly available on Sourceforge.net,[3] one of the most popular repositories of open source software offering online resources for open source software development and content creation. Since its first beta release in September 2006, CoreASM has gone through a number of revisions and its latest version (under testing at the time of writing this document) offers substantial improvements over its previous version in terms of both features and performance.

The rest of this chapter continues with an overview of the architecture of CoreASM in Java. Section 7.2 looks into the implementation of the CoreASM engine focusing on the implementation of the two more challenging components, the Abstract Storage and the Parser, and the implementation of CoreASM plugins. Section 7.3 concludes this chapter by introducing the tools and user interfaces that are built around the CoreASM engine.

7.1 The Architecture

The CoreASM engine has a micro-kernel architecture. Recalling the architecture of CoreASM as presented in Chapter 4, the kernel of the engine provides only the essential aspects of the engine required for the plugins and applications to be built upon. Furthermore, the kernel is decomposed into four components: a parser, an interpreter, an abstract storage, and a scheduler. The interface of the engine to its environment (and in parts, to its four components) is provided by a special component called the Control API (see Figure 4.2).

Closely following the design of the engine, the Java implementation of CoreASM implements the kernel of the engine in terms of four components and a Control API. The interface of the components are defined by four Java interface files: `Parser`, `Interpreter`, `AbstractStorage`, and `Scheduler`. For every component, a default implementation is provided in form of a Java class file. However, every component is carefully encapsulated in its interface and, as a result, a different implementation can be used as long as it complies with the the interface of the component and its specification. Since Control API acts as a double interface, providing services both to the environment of the engine and to its internal components—the former being a subset of the latter, two Java interface files together define the interface of the engine: (i) a `CoreASMEngine` interface defines the interface of the engine to its outside environment offering services such as loading, parsing, or execution

[3] `http://www.sourceforge.net`

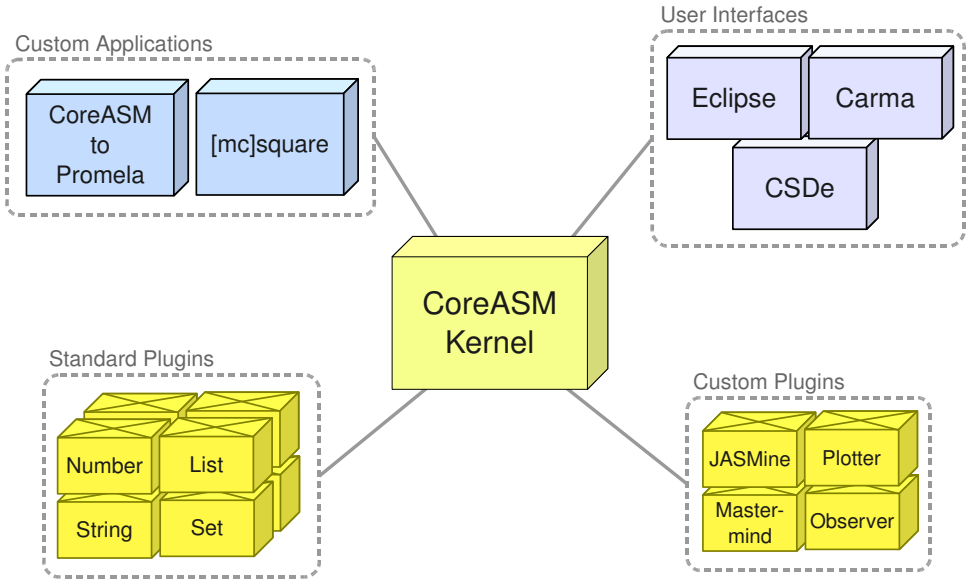

Figure 7.1: CoreASM Kernel, Plugins, and Applications

of specifications; (ii) a `ControlAPI` which extends the `CoreASMEngine` interface providing access to every component, a mapping of plugin names to actual plugin instances, and error reporting services. An implementation of the CoreASM engine is provided by the Java class file `Engine` which implements the `ControlAPI` interface.

The `CoreASMEngine` interface provides a comprehensive interface to the engine. Through this interface, applications can (i) load CoreASM specifications into the engine, execute them step by step, and access the simulated state and the latest update set throughout the execution, (ii) use the engine as a parser to just parse specifications into parse-trees (which can then be externally processed for various purposes such as model checking [59, 93]), or access the list of plugins required by a given specification, (iii) modify various engine properties and also observe the behavior of the engine by implementing the `EngineObserver` interface.

There are currently two user interfaces available for CoreASM (see Figure 7.1): a comprehensive command-line user interface, called Carma, and a graphical interactive development environment in the Eclipse platform, known as the CoreASM Eclipse Plugin. There is also a sophisticated tool under development for creating and modifying Control State ASMs and translating them into CoreASM specifications, called CSDe. Section 7.3 presents these tools

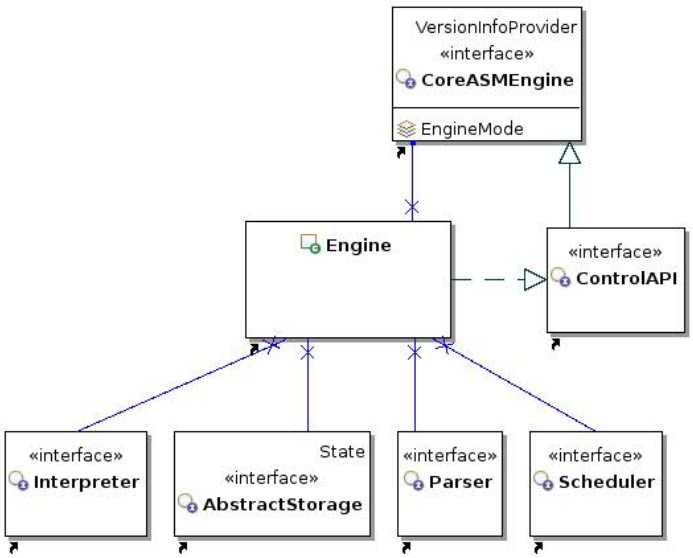

Figure 7.2: Components of the CoreASM Engine

in more detail.

The CoreASM kernel also defines the skeleton of a CoreASM plugin in form of a Java abstract class `Plugin`. Various types of extensions that plugins can provide to the engine, such as parser extension or vocabulary extension (see Section 5.5 for a complete list), are defined in terms of Java interface files. Every CoreASM plugin must extend the `Plugin` abstract class and most likely implement one or more of the extension interfaces to offer its contribution to the engine.

7.2 The CoreASM Engine

In this section we briefly look into the implementation of the kernel (focusing on the more challenging components, the Abstract Storage and the Parser) and the plugin framework.

7.2.1 The Kernel

CoreASM engine is represented by the `CoreASMEngine` interface and is implemented by the `Engine` class file which serves two purposes: (i) it provides an implementation for the interface of the engine to its outside environment, and (ii) it acts as a container for the main components of the engine and maintains the control state of the CoreASM engine. In

order for the engine to be always responsive to its environment, the `Engine` object runs in two parallel processing threads: one, being the environment or the *caller*'s thread, responds to requests from the environment (such as sending commands, setting engine properties, or retrieving updates) and the other one maintains the internal control flow of the engine.

The Abstract Storage

The Abstract Storage is implemented by more than three dozen classes in the package `org.coreasm.engine.absstorage`. A hierarchy of classes implement various types of elements defined in the kernel (see Figure 7.3). At the root of this hierarchy, we have the `Element` class which is the superclass of all the values in CoreASM states, implementing the ELEMENT domain. Following the specification of Section 5.1, every instance of `Element` has a background and a notion of equality. Three immediate subclasses `BooleanElement`, `RuleElement`, and `FunctionElement` respectively implement the domains of BOOLEANELEMENT, RULE, and FUNCTIONELEMENT defined in Section 5.1. The domain of BACKGROUNDELEMENT and UNIVERSEELEMENT are implemented by similarly named subclasses of a more generic class `AbstractUniverse` which captures similar aspects of these two domains. Since only a finite set of elements can be represented by Universe elements, `UniverseElement` also implements the `Enumerable` interface.

The main class of this package is `HashStorage`, which offers an implementation for the Java interface `AbstractStorage` based on hash tables. The CoreASM state is implemented by the Java class `HashState` through three separate mappings of names (Java `String` values) to Function elements (instances of `FunctionElement`), Rule elements (instances of `RuleElement`), and Background and Universe elements (instance of `AbstractUniverse`), thereby implementing contents of CoreASM state as defined in Section 5.1:

$$\begin{aligned} &\mathit{stateFunction} : \textsc{State} \times \textsc{Name} \mapsto \textsc{FunctionElement} \\ &\mathit{stateRule} : \textsc{State} \times \textsc{Name} \mapsto \textsc{Rule} \\ &\mathit{stateUniverse} : \textsc{State} \times \textsc{Name} \mapsto \textsc{UniverseElement} \end{aligned}$$

The Parser

Implementing the parser component of the CoreASM engine was quite a challenge. At first, we were looking for fast and efficient parser generators that can be called upon loading a specification to generate a parser based on the grammar provided by the specific plugins

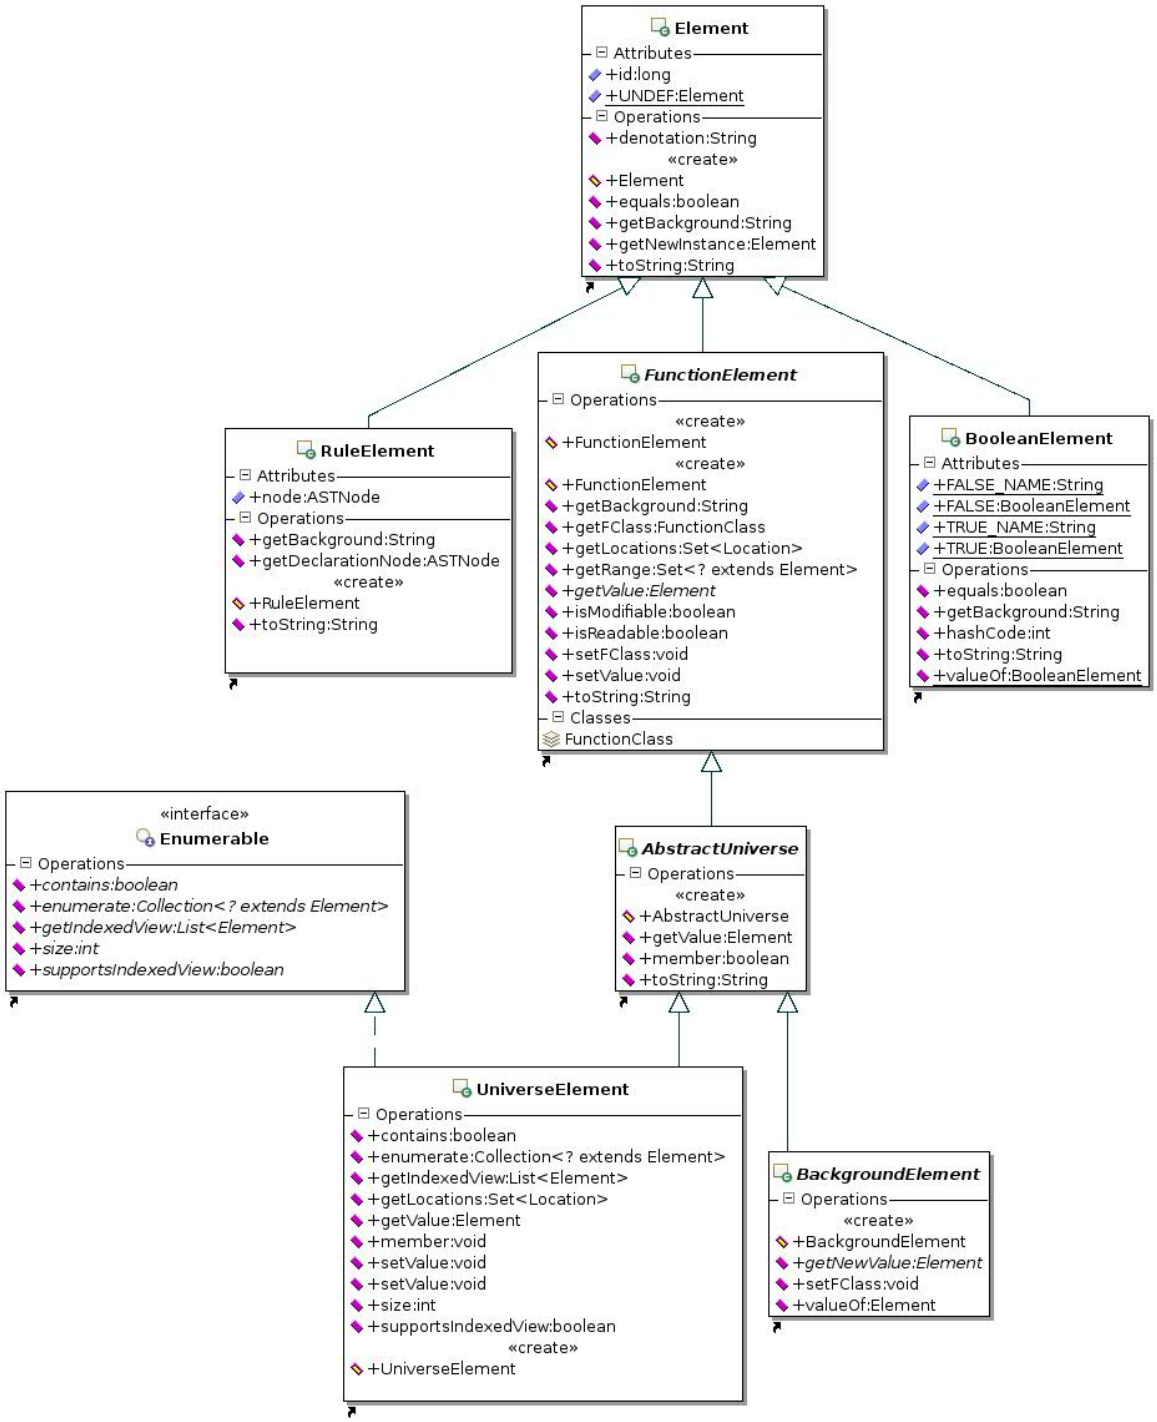

Figure 7.3: Core Elements Defined in the Abstract Storage

that are used in that specification. Originally, we used the OOPS (Object Oriented Parser System) parser generator[4] developed and maintained by Axel-Tobias Schreiner and his students Bernd Kühl and William Leiserson. The original OOPS parser generator was quite restrictive for CoreASM as it would generate only LL(1) parsers. Later, Will Leiserson extended and improved OOPS into an LL(k) parser generator [90]. However, the new parser generator was not fast enough on typical CoreASM specifications to be used every time a specification is being loaded.

We looked into a number of available open source parser generators in search of an efficient LL(k) parser generator written in Java and we eventually found jparsec,[5] a recursive-descent parser combinator framework written for Java. In contrast to traditional parser generators like YACC or ANTLR, jparsec grammar is written in native Java language and is defined in terms of special Java instances of a `Parser` class. Each parser object represents a grammar rule and can be combined with other parser objects to create more complex production rules. For example, a production rule of the form "$A ::= B \mid C \mid D$" can be created by the following Java code:

```
Parser<Foo> a = Parsers.or(b, c, d);
```

where `b`, `c`, and `d` are parser instances representing the non-terminals B, C, and D in our production rule. In jparsec, once a parser object is created, it can be asked to "parse" a given input:

```
a.parse("text to be parsed");
```

Depending on how the parsers are defined, the return value (the result of parsing) can be a value resulting from a calculation or an abstract syntax tree representing the input text.

This feature of jparsec appeared to be very beneficial for CoreASM. Upon loading a specification, the kernel provides references to the core parser objects (such as white spaces, identifiers, terms, etc.)[6] and make them available for plugins to build upon. Plugins in turn provide their contributions to the parser in form of new jparsec parser objects. The kernel then puts all these contributions together to create the final parser that will be used to parse the specification.

[4] http://www.cs.rit.edu/~ats/

[5] http://jparsec.codehaus.org

[6] Some of these core parsers, such as the one for parsing CoreASM terms, can also be extended by plugins.

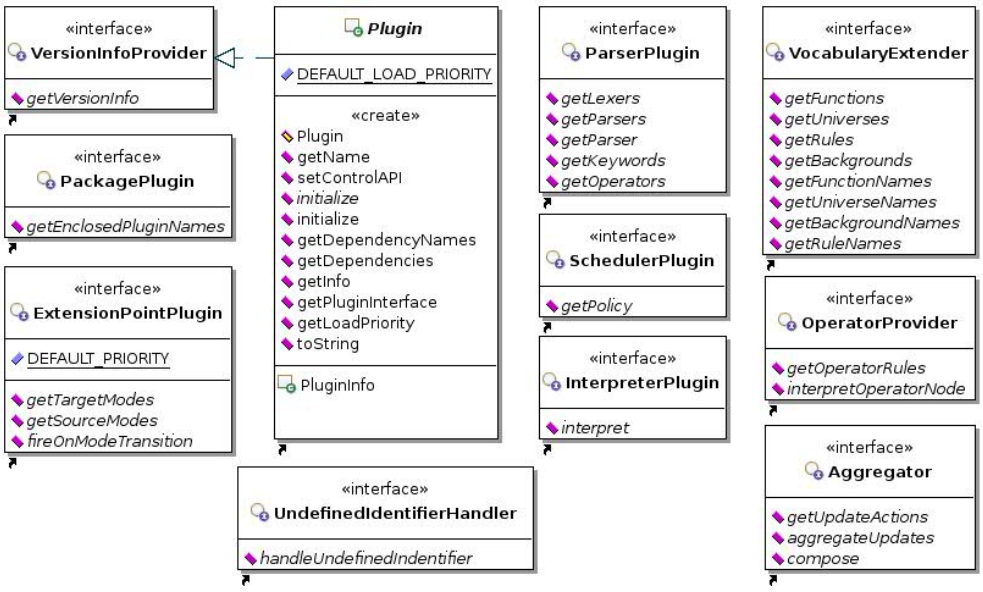

Figure 7.4: CoreASM Plugin Interfaces

7.2.2 CoreASM Plugins

Every CoreASM plugin must extend the abstract class `Plugin` and most likely implements at least one of the nine plugin interfaces offered by the engine (see Figure 7.4).[7] We introduced the seven most important plugin interfaces in Section 5.5; the remaining two are the `PackagePlugin` and the `UndefinedIdentifierHandler` interface. The former should be implemented by plugins that are defined to serve as a "package" of other plugins. For example, CoreASM comes with a *Standard Plugin* which is a plugin that implements only the `PackagePlugin` interface and when loaded (see LoadSpecPlugins on page 52) provides a list of plugins that it consists of. The latter one, `UndefinedIdentifierHandler`, is implemented by plugins that offer a mechanism to deal with undefined identifiers. For example, a plugin can implement this interface and override the default behavior of the engine and generate an error whenever an undefined identifier is recognized by the engine; see Section 5.2.2 and the definition of the rule HandleUndefinedIdentifier in Section A.2.

A CoreASM plugin is most likely accompanied by a number of auxiliary Java classes. As a result, every CoreASM plugin is expected to be packed into a single JAR file[8] together

[7] Even if a plugin does not implement any of the plugin interfaces, it is still a valid plugin as long as it properly extends the `Plugin` class. However, the effect of loading such a plugin would be extremely limited.

[8] JAR (Java Archive) files are package files that are used by software developers to distribute Java classes

with an identification file. When an instance of `Engine` is initialized, it searches a specific *plugin* folder, creates a catalog of available plugins (abstractly modeled by the LoadCatalog rule on page 50) and loads the plugin class files together with their corresponding classes into the Java Virtual Machine (JVM), so that they can be later instantiated if needed. As a result, to add a new plugin to CoreASM, one only needs to put the JAR file of the compiled plugin into the *plugin* folder of the engine.

7.3 User Interfaces and Tools

The CoreASM engine is implemented as a Java component and requires a *driver* program (such as a user interface) to run the engine, e.g., to pass specification files to the engine and to control its simulation run by manipulating parameters. There are currently two user interfaces available for the CoreASM engine: a powerful command-line tool called Carma, and a graphical interactive development environment in the Eclipse platform, known as the CoreASM Eclipse Plugin.

Carma

Carma is a comprehensive command-line user interface for CoreASM that offers rich control over the runs of the engine through more than a dozen command-line options and switches. To execute a specification, users can simply run Carma on the command line and pass it the name of the specification file as an argument. By default, Carma does not have a termination condition, but it offers a number of termination conditions, such as termination after a number of steps, termination on empty updates, and termination when there is no valid agent with a defined program. As an example, the following command runs the CoreASM specification `MySpec.coreasm` using Carma and stops after 30 steps or after a step that generates empty updates; it also provides a print-out of the final state before termination.

```
carma --steps 30 --empty-updates --dump-final-state MySpec.coreasm
```

---

and their associated metadata.

The CoreASM Eclipse Plugin

The CoreASM Eclipse Plugin is a graphical interactive development environment for CoreASM in form of a plugin for the well-known Eclipse software development platform. The IDE provides various options to control execution of CoreASM specifications. The plugin extends the Eclipse platform to support dynamic syntax highlighting and interactive execution of CoreASM specifications. Since the language of CoreASM for a given specification is defined by the set of plugins used by that specification, with every change to the specification, the editor component of the CoreASM Eclipse Plugin passes the specification to the CoreASM engine and gets the set of plugins that are used by the specification. The editor then asks the plugins for the set of keywords, functions, universes and backgrounds they provide and uses this information to offer a dynamic syntax highlighting of the specification.

Figure 7.5(a) shows a snapshot of the CoreASM environment in Eclipse. At the top left corner (1), the toolbar is extended to include buttons to pause, resume and stop a simulation run. The editor (2) provides dynamic syntax highlighting for CoreASM specifications based on the set of CoreASM plugins used in the specification. A configurable output console (3) provides a print-out of the results of the simulation with optional additional information on the simulation process and the state of the simulated machine.

7.3.1 CSDe

The Control State Diagram editor (CSDe), under development by Piper J. Jackson [51], is a sophisticated tool for creating and modifying Control State ASMs and translating them into CoreASM specifications. The tool is implemented as a plugin for the Eclipse software development platform. The plugin allows the user to work with Control State Diagrams (CSDs) using a point-and-click schema (see Figure 7.5(b)).

The simplicity of control state diagrams and the intuitiveness of the graphical user interface work together to allow users to confidently contribute to the design, regardless of their technical background. The diagram editor (CSDe) is capable of automatically transforming diagrams into CoreASM specifications. Since control state diagrams do not necessarily include initial states of the system or other more concrete information required for machine execution, such specifications may not be directly executable. However, they provide an abstract structure for the design of systems and act as foundations for further development

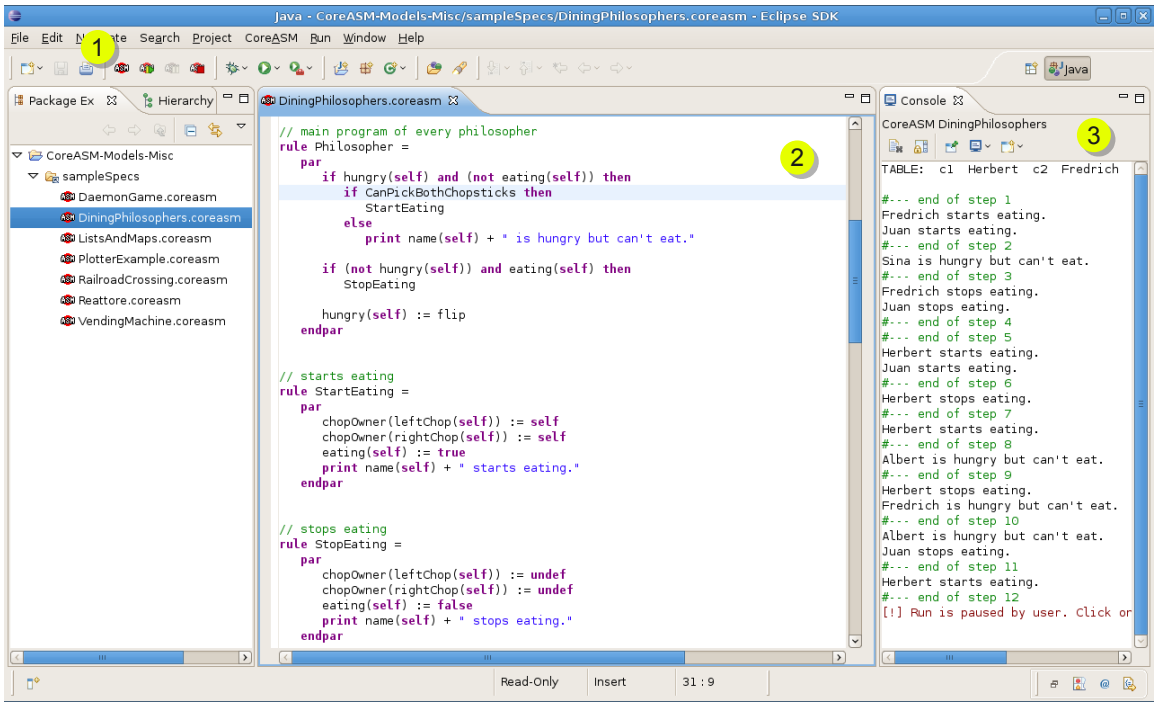

(a) CoreASM Eclipse Plugin

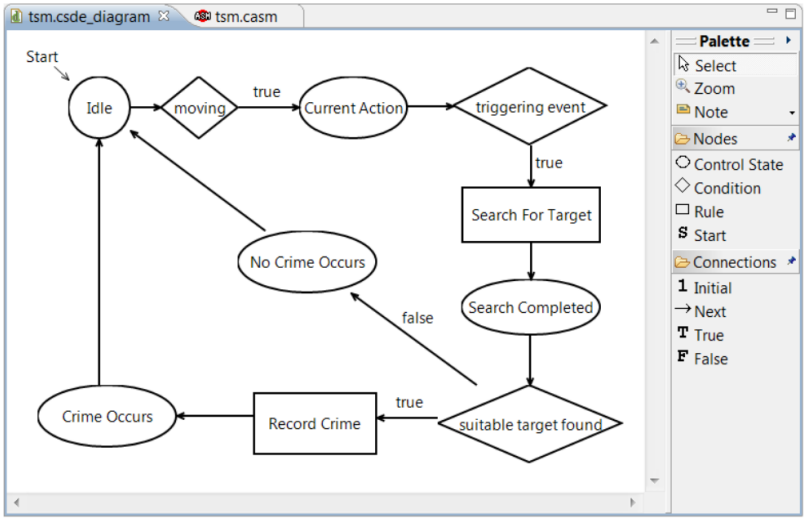

(b) CSDe: A Control State Diagram editor for CoreASM

Figure 7.5: CoreASM Tools in Eclipse

of the specifications. The automatic translation feature facilitates the transition from high-level design ideas expressed in graphical form towards abstract yet relatively more concrete specifications.

7.3.2 Model Checking CoreASM Specifications

The CoreASM engine facilitates experimental validation of ASM models by providing the means to execute abstract specifications and to explore behavioral aspects in an interactive fashion. However, experimental validation without model checking cannot formally verify the correctness of a system with respect to all of its possible behaviors. In order to provide model checking support for CoreASM, George Ma developed a tool called CoreASM2Promela [93] that utilizes the CoreASM engine to translate CoreASM models into equivalent Promela models which can be verified using the Spin model checker.[9] From a high level perspective, the steps in the translation and verification process are as follows: (i) a CoreASM specification is loaded and parsed by the CoreASM engine, producing an abstract syntax tree; (ii) the tree is translated into Promela; (iii) Spin is invoked to generate a verifier of the Promela model, producing C code; (iv) the C code is compiled, generating a custom verifier of the CoreASM specification; (v) the verifier is run, producing a counter example if the property being checked does not hold.

In order to properly translate CoreASM specifications into Promela models, we needed to extend the CoreASM language by two new plugins, namely the *Signature Plugin* (see Section 6.4.1) and the *Property Plugin*, to support declaration of function signatures and specification of LTL properties as part of CoreASM specifications. The Property Plugin is a small plugin that allows correctness properties, expressed as LTL formulas, to be included in the header of a CoreASM specification. Presently, specified properties do not have any meaning during ASM simulations (although it may be possible to extend the Property plugin to check simple global assertions). Correctness properties are only applicable during model checking, and are translated by our CoreASM to Promela translator.

The Property plugin provides the following pattern to declare new LTL properties:

[**check**] **property** *LTL-property*

[9] Spin is a widely used automata based model checker that has been used extensively in the design of asynchronous distributed systems [80].

Including the keyword **check** with a property declaration indicates that the property should be checked during model checking.

Since Spin does not allow LTL properties to be included directly in a specification, the Property plugin was developed to improve the usability of the model checker. In Spin, properties are defined by describing the behavior of a property automaton. Moreover, Spin only allows a single property automaton in each model, while the Property plugin allows multiple properties to be specified for a single specification.

George Ma has successfully used CoreASM2Promela to model check several non-trivial ASM specifications; the details of the case studies and a comprehensive discussion of the results are presented in George Ma's M.Sc thesis [93]. However, there are certain limitations in model checking abstract ASM specifications using Spin. For example, as Spin can only check finite models, the translation scheme is limited to CoreASM specifications which have finite states. Thus, the translation supports only static universes and enumerated backgrounds.

Chapter 8

Case Studies

This chapter presents three case studies from three diverse application contexts to examine the practicability of using CoreASM for requirements analysis, design specification and rapid prototyping of abstract system models. These three examples result from projects that have been carried out at SFU's Software Technology Lab[1] in close collaboration with industrial partners, Defence R&D Canada, SFU's Institute for Canadian Urban Research Studies and the Royal Canadian Mounted Police,[2] and they illustrate the wide scope of application domains for CoreASM, beyond classical software system design problems.

8.1 The DRCMA Project

Dynamic Resource Configuration & Management Architecture (DRCMA) [58, 52] is a highly adaptive and auto-configurable multi-layer network architecture for distributed information fusion. The primary goal of DRCMA is to address large volume surveillance challenges, assuming a wide range of different sensor types operating on multiple mobile platforms for intelligence, surveillance and reconnaissance. The focus is on network enabled operations to efficiently manage and improve employment of heterogeneous sets of surveillance and patrolling resources, their information fusion engines and their networking capabilities under dynamically changing and essentially unpredictable conditions. The architecture is built on realistic application scenarios adopted from the design and development of the

[1] http://stl.sfu.ca

[2] This chapter provides a summary of our previously published materials on these subjects [51, 58, 52, 50].

CanCoastWatch system [123, 57].

8.1.1 Objectives and Challenges

The overall design objective of DRCMA is a highly robust and scalable network architecture that supports reconfigurable applications and self-organizing structures, flexibly adapting to dynamically changing resource requirements as well as changes in the availability of resources. Global mission goals are to be operationalized into local tasks performed by semi-autonomously operating resource units that can handle basic adjustments and realignments of resources automatically. The architectural design emphasizes a hierarchical command and control structure.

Missions injected into the system are complex tasks, each of which needs to be transformed into a collection of constituent elementary tasks, so as to map these tasks onto the available resources. Complex tasks therefore are decomposed in one or more steps into simpler ones until all of the resulting tasks are of elementary type, meaning that each of them can directly be assigned to a *physical resource* capable of performing the task. Physical resources refer to individual resource entities that exist in the physical environment. Depending on the level of abstraction, a physical resource may either identify a group of sensor platforms or a single sensor platform or even an individual sensor unit on a sensor platform.

Logical resources represent clusters of resources formed by aggregating two or more physical and/or logical resources, each with a certain range of capabilities, into a higher level resource with a greater capacity for performing complex operations. Resource clusters operate semi-autonomously to increase robustness and to reduce control and communication overhead by making local decisions regarding the realignment and reorganization of resources within the cluster. Dynamic reconfiguration of clusters is performed in an ad hoc manner using 'plug and play' mechanisms. Resources may join or be removed from a cluster on demand and depending on their capabilities, geographic location, cost aspects and other characteristics.

8.1.2 Conceptual Model

The overall organization of the dynamic resource configuration and management architecture assumes a hierarchical command and control structure that is described by a hierarchical

network of nodes representing mobile resources as illustrated in Figure 1. There are two basically different types of resource entities as follows:

- *Physical resources* refer to real-world resource entities as part of an existing distributed fusion system. In the hierarchical structure, only the leaf nodes represent physical resources. Depending on the level of abstraction at which a distributed fusion system is considered, a physical resource may refer to a group of mobile sensor platforms, to a single mobile platform or to an individual sensor unit on a sensor platform.
- *Logical resources* refer to abstract resource entities formed by clustering two or more physical and/or logical resources, each with a certain range of capabilities, into a higher level resource with aggregated (richer) capabilities needed to perform complex operations. A logical resource identifies a *cluster of resources*[3]. In the hierarchical structure (e.g. as illustrated in Fig. 8.1(a)) all non-leaf nodes represent logical resources.

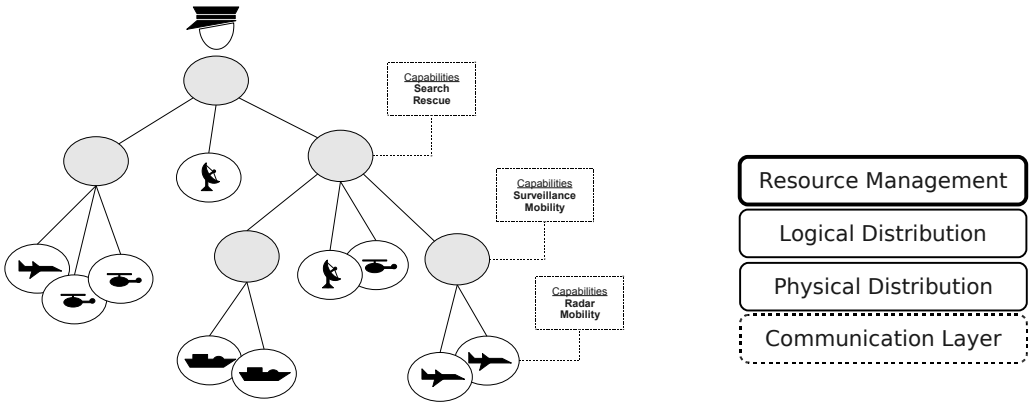

(a) Nodes Represent Physical or Logical Resources. (b) DRCMA Service Layers

Figure 8.1: Architectural View of DRCMA

Resource clusters operate semi-autonomously to increase robustness and to reduce control and communication overhead by making local decisions regarding the realignment and reorganization of resources within the cluster identified by a logical resource. Dynamic reconfiguration of clusters is performed in an ad hoc manner using 'plug and play' mechanisms. Resources may join or be removed from a cluster on demand and depending on

[3]Logical resources require *command* capabilities in their clusters, but we abstract away from this notion in our model.

their sensor capabilities, mobility capabilities, geographic location, cost aspects and other characteristics. The underlying design principles resemble those for improving performance and robustness in mobile ad hoc networks [96].

Resource Allocation

Missions represent complex tasks, typically involving a number of explicitly or implicitly identified subtasks, each with specific resource capability requirements that need to be matched with the capabilities of mobile resources in order to perform any such task. In general, the operations carried out by a set of resources allocated to a mission are distributed in time and space, making the coordination of these resources a challenging problem.

A mission is decomposed into subtasks in one or more steps until all of the resulting tasks are of elementary type, meaning that each of them can directly be assigned to a physical resource capable of performing the task. New missions are introduced by assigning the mission to the top-level node. When a logical resource receives a complex task, it tries to find a child node with matching capabilities to perform the task. If none of the child nodes can perform the task, the logical resource attempts to split the task into a collection of subtasks that can be performed by two or more of its child nodes. If this attempt fails as well, the task will be rejected.

Intuitively, one may view new tasks as 'sinking' downwards in the node hierarchy until they reach a matching physical resource or become transformed into a collection of related subtasks.

Resource Distribution

The distribution of mobile resources can change dynamically over time depending on a number of internal and external factors. To clearly separate concerns, DRCMA distinguishes two types of resource distribution:

- *Physical distribution* refers to the spatio-temporal distribution of mobile resources in the geographical environment. Position information and projections of resource trajectories provide important input to support the logical distribution (e.g., keeping resources of the same group in close proximity to each other) and also to satisfy communication requirements (e.g., moving a resource to a location in order to act as a communication proxy).

- *Logical distribution* refers to the dynamic configuration of physical resources into clusters that change in response to the changes in tasking information (e.g., new task orders or changes in task priorities), changes in the capabilities of resources (e.g., device failures or new resources joining the network), changes in the environment (e.g., changes in weather conditions), and to maintain a desirable workload balance within individual clusters and across the whole network.

Layered Architecture

Dynamic resource management policies govern the migration of resources between clusters based on common prioritization schemes for resource selection, load balancing and organization of idle resource pools. Configuration and management of resources across nodes is organized by building on a service-oriented architecture consisting of four hierarchical *service layers*, namely: Resource Management (L_4), Logical Distribution (L_3), Physical Distribution (L_2), and Communication (L_1), where L_1 refers to the bottom level layer; see Figure 8.1(b).

The proposed DRCMA model assumes clearly identified and well defined interfaces between layers, such that $layer_n$ renders services to the next higher $layer_{n+1}$ using $layer_n$ protocols realized by means of services provided by $layer_{n-1}$ [70]. The encapsulation of services in separate layers not only enhances a clear separation of concerns but also simplifies the control of complexity by providing convenient abstractions for decomposing complex interaction patterns.

8.1.3 Formal DRCMA Model

Starting from the design concepts described above, we turn the abstract DRCMA view into a high-level formal model [58, 52] that can be systematically analyzed, inspected, refined and validated. The formal representation ensures that the key system attributes are specified concisely and unambiguously, providing a reliable foundation for checking that these attributes are well understood and actually do meet the functional requirements.

Nodes are the basic components of the DRCMA model. Every node refers to either a logical or a physical resource in the network (see Figure 8.1(a)). Every logical resource represents a non-empty set of subordinate resources, called *child resources*, that belong to its cluster. Hence, we have:

Resources

universe NODE
universe RESOURCE

$node$: RESOURCE $\mapsto$ NODE
$resource$: NODE $\mapsto$ RESOURCE
$\forall r \in \text{RESOURCE}\ \forall n \in \text{NODE}\ \ resource(n) = r \Leftrightarrow node(r) = n$

$cluster$: RESOURCE $\mapsto$ SET(RESOURCE)
$isCluster$: RESOURCE $\mapsto$ BOOLEAN
$\forall r \in \text{RESOURCE}\ \neg isCluster(r) \Rightarrow (cluster(r) = \{\})$

Resources may have different kinds of capabilities; for instance, a helicopter can have 'mobility', 'communication', and 'radar'. Every capability has a number of attributes describing it in more detail. An *attribute* is defined as a pair $\langle name,\ value \rangle$. For instance, a 'mobility' capability may have a 'speed' attribute with a numeric value such as $\langle speed,\ 50kmph \rangle$. One may specify additional aspects of mobility and define other attributes such as 'type', taking values from a given set {'air', 'ground', 'water'} to more precisely describe the capability.

Resources can have more than one instance of a capability. For instance, the helicopter in our example may be able to communicate via two different communication links. Therefore, it has two 'communication' capabilities, each with a different value for its 'type' attribute: $\langle type,\ link11 \rangle$, and $\langle type,\ link16 \rangle$. We can formally define capabilities and their relationship to resources as follows:

Capabilities

domain CAPABILITY

$capabilityName$: CAPABILITY $\mapsto$ NAME
$capabilityAttribute$: CAPABILITY $\mapsto$ NAME $\times$ VALUE

$nodeCapabilities$: NODE $\mapsto$ MULTISET(CAPABILITY)

Moreover, each node has a multiset of capabilities[4] consisting of all the capabilities it can provide. For a higher level node (logical node), which itself consists of other physical or logical nodes, the multiset represents the aggregation of all the capabilities provided by

[4]Since a node may have more than one instance of a specific capability, we model the collection of capabilities as a multiset.

its children.

Different capabilities may be required to perform a single task. In addition, various capabilities may each satisfy a single capability requirement of a task. For instance, two different resources, SAR helicopter and patrol aircraft, both having a *mobility* capability of type 'air', may each satisfy the capability requirement of a task that requires air mobility. In order to model the matching of capability requirements to capabilities, we model capability requirements as *Capability Patterns* with almost the same structure that capabilities have, with the conceptual difference that the attribute values of a capability pattern can be a range or a set of acceptable values. A *matchCapability* function of the form

$$matchCapability : \textsc{CapabilityPattern} \times \textsc{Capability} \mapsto \textsc{Boolean}$$

holds if a certain capability matches a given capability pattern.

Capability Pattern

domain CAPABILITYPATTERN

requiredCapabilities : TASK $\mapsto$ SET(CAPABILITYPATTERN)
decomposeCapability : CAPABILITYPATTERN $\mapsto$ SET(CAPABILITYPATTERN)

Hence, *requiredCapabilities*(t) is a set of *Capability Patterns*, describing the capabilities a task t needs in order to be performed. Capability requirements of a task can be satisfied by finding a matching capability for each of its capability patterns.

Logical Distribution The logical distribution of resources is reflected by the network topology as stated through the following functions defined on nodes:

- *rootNode* : $\mapsto$ NODE, a static function that identifies a distinguished node of the network representing the top level command and control unit.
- *childNodes* : NODE $\mapsto$ NODE-SET, *childNodes*(*node*) holds the set of nodes under direct authority of *node*.
- *parentNode* : NODE $\mapsto$ NODE, points to the parent node of a node.

Physical Distribution The physical distribution of resources within a given geographical environment in which they operate typically changes over time. This is abstractly modeled by a monitored function

$$\textit{location} : \textsc{Resource} \mapsto \textsc{Coordinate}$$

which, in any given system state, associates with each of the resources a geographical location, e.g. as identified by a global positioning system. As resources change their location dynamically, the function *location* may change its interpretation from state to state. Based on the location of resources and certain dynamic characteristics of the environment, such as weather and terrain conditions, various derived functions model various aspects of the physical distribution of resources. For instance, for any given node, the set of all the given nodes that are reachable over a certain communication channel is modelled by

$$\textit{visibleResources} : \textsc{Resource} \times \textsc{Channel} \mapsto \textsc{Resource-set}.$$

The physical distribution layer (see Figure 8.1(b)) provides services to query about and manipulate the physical distribution of resources in DRCMA. We abstract here from the internals of this layer, assuming an underlying model resembling those used in the routing layer of mobile ad hoc communication networks [70].

Dynamic Resource Management

DRCMA handles task allocations and maintains a dynamic logical distribution of resources by actively monitoring the operational status of resources and the tasks to be performed, reacting to changing situations. The following events can trigger a change in the logical distribution of resources: *1)* occurrence of new or modified tasks; *2)* changes in resource capabilities or availability of resources (e.g., failure of sensor units); and *3)* problems with communication links.

In response to an event that triggers a reaction, the network configuration can be changed by applying the following canonical *transformation patterns* (see also Figure 8.2):

1. New clusters can be created on demand, e.g., in response to a new high priority task.
2. Resources can be moved from one cluster to another, e.g., in order to better serve a higher priority mission or to balance resource load.
3. Resources can form their own solitary clusters if they get disconnected (and isolated) from other resources due to a communication failure.

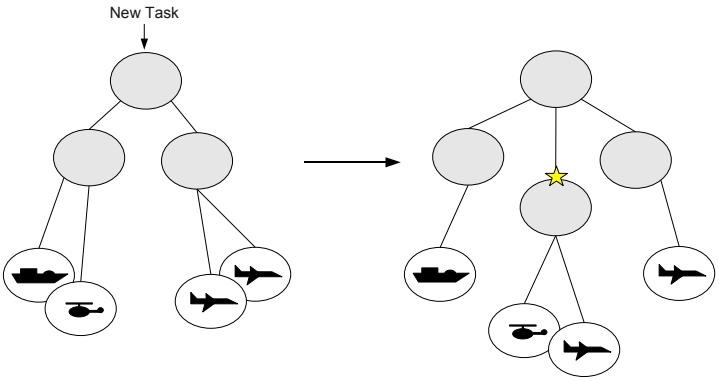

(a) New clusters can be created based on new tasks received.

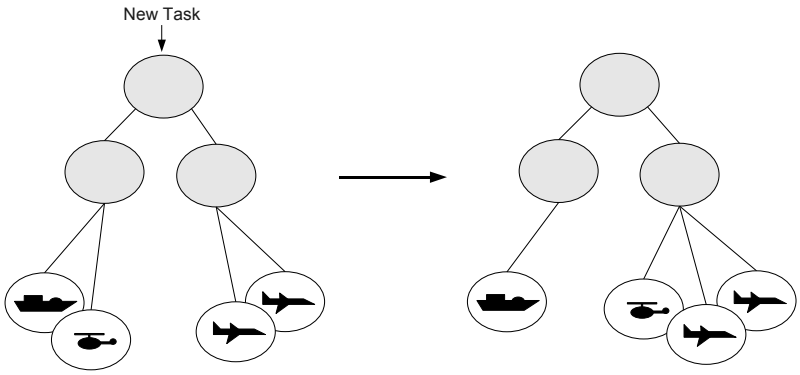

(b) Resources can be moved from one cluster to another in response to new tasks.

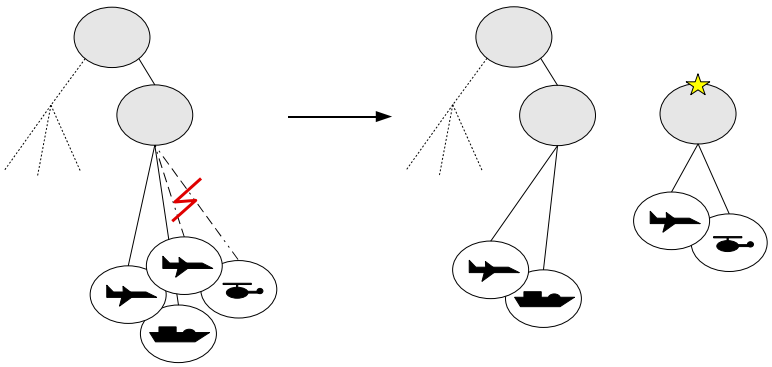

(c) Resources can form their own solitary clusters if disconnected.

Figure 8.2: Basic Transformation Patterns

4. Clusters can be merged into a larger cluster to satisfy a new or changed goal. This can be considered as a special case of 1.

The process of maintaining the configuration of resources is modeled as a distributed process performed by the individual nodes of the network. Every node actively monitors the environment and maintains the configuration of its resources to maintain stability, balance workload, and increase resource performance. The behavior of the DRCMA nodes is captured in four layers (see Figure 8.1(b)):

DRCMANodeProgram ≡
ResourceManagementProgram
LogicalDistributionProgram
PhysicalDistributionProgram
CommunicationProgram

The behaviour of DRCMA resource management layer is decomposed into five main activities running in parallel, as stated by the following ASM program formed by the parallel composition of four behavioral component descriptions.

ResourceManagementProgram ≡
MonitorNewTasks
MonitorResources
MonitorComLinks
ProcessObservedEvents$_{RM}$
RespondToMessages$_{RM}$

Here, as an example, we look into the refinement of MonitorNewTasks and specify how DRCMA nodes manage new task assignments.

8.1.4 New Task Assignments

When a new task t is assigned to a node N, either by a parent node or as a result of a situation analysis performed by the node itself, N uses a process similar to a 'call for tender' to find a suitable set of resources that together can handle task t (see also Figures 8.2(a) and 8.2(b)):

1. Compute the capabilities required for processing task t.

2. Ask all the child nodes of N if they could provide these capabilities and, if so, what is the estimated cost of providing such capabilities.

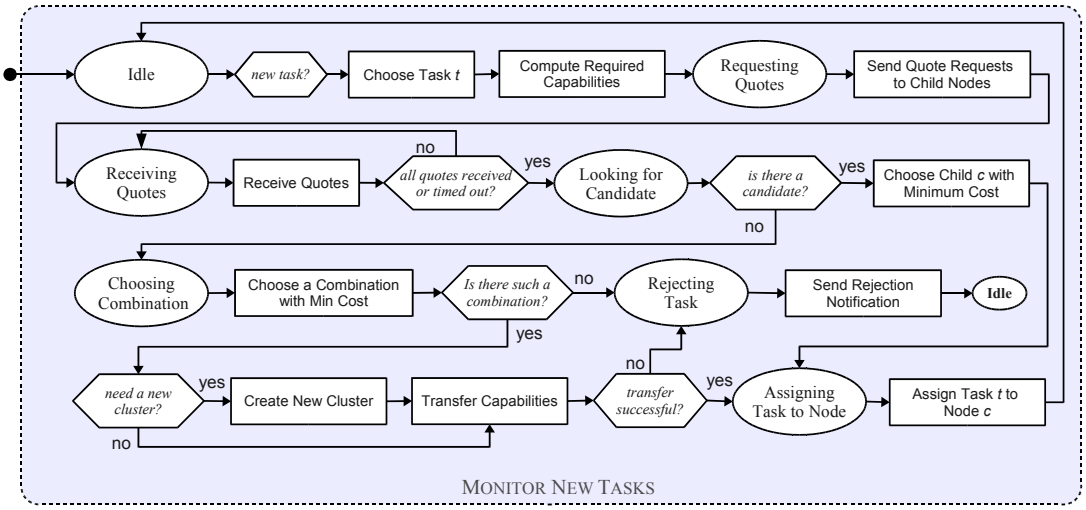

Figure 8.3: Control State ASM of Monitoring New Tasks by Logical Nodes

3. If there is at least one potential candidate, i.e. a child node that can provide all the required capabilities, choose one such candidate c with the minimum total cost, and assign task t to c.

4. If there is no such child node, choose a combination of child nodes that together can perform task t with the minimum total cost.

   (a) Decide whether a new cluster should be created or a currently available one can be used; call that cluster nc.

   (b) Send a logical distribution change request to move the required capabilities from other nodes to cluster nc, and wait for receiving the acknowledgement from all the nodes in the sub-cluster.

   (c) If all the required resources (capabilities) are transfered properly under nc, assign task t to nc.

5. If no combination of nodes can be found to handle task t, reject task t and notify the parent node even if task t is emerged from within node N itself.

As part of our formal model, we provide a control state ASM [25] diagram formally defining the process of monitoring new tasks (see Figure 8.3), in which ellipses represent

control states, diamonds represent conditionals, and rectangles represent ASM rules (i.e., actions). While the diagram provides the overview of the process and its control flow, the details of the conditions and the actions they trigger are defined in terms of an ASM. For instance, the following predicates specify the condition '*Is there a candidate?*':

$$candidateExists(node, task) \equiv \exists\ c \in childNodes(node)\quad isACandidate(c, task)$$

$$isACandidate(node, task) \equiv \forall\ rc \in requiredCapabilities(task)\quad quote(node, task)(rc) \neq undef$$

and the following rules define the actions 'Send Quote Requests to Child Nodes' and 'Receive Quotes' of the resource management (RM) layer:[5]

SendQuoteRequests$_{RM}(task) \equiv$
  forall c **in** $activeChildNodes(self)$ **do**
    let $m = new(\text{MESSAGE})$ **in**
      $msgType(m) := quoteRequest$
      $msgData(m, \text{"task"}) := task$
      $msgData(m, \text{"capabilities"}) := requiredCapabilities(task)$
      $quote(c, task) := undef$
      $\text{SendMessage}_{RM}(c, m)$
where
  $activeChildNodes(n) \equiv \{x \mid x \in childNodes(n) \wedge \neg isDead(x)\}$

ReceiveQuotes$_{RM}(task) \equiv$
  choose $m \in inbox(self)$ **with**
      $msgType(m) = quoteResponse \wedge msgData(m, \text{"task"}) = task$ **do**
    $quote(sender(m), task) := msgData(m, \text{"quote"})$
    remove m **from** $inbox(self)$

The above rules exhibit two important abstractions in the DRCMA model. In these rules, the abstract view of communication services and data structures of messages allows us to focus on the main functionality of the process. To send quote requests to child nodes, DRCM relies on the messaging services provided by the communication layer. For every child node, a new *quote request* message is created asking for a quote on the cost of performing the new task, and the message is sent using the abstract routine SendMessage_{RM}. In the

[5] In these rules, *self* refers to the node executing the rule.

next step, the node looks into its message inbox and non-deterministically chooses quote messages related to the new task and stores the received quotes in an internal data structure to be used later in the process (see ReceiveQuotes above). This rule is repeated until all the expected quotes have either been received or timed-out.

If no single resource cluster suitable to perform the new task is available, the best combination of resources from different clusters are selected to form a new cluster (see "Choosing Combination" in Figure 8.3). To create a new cluster, a new logical resource and a corresponding node is created, and the hierarchical structure is then modified by changing the values of functions *parentNode* and *childNodes*, effectively adding the new cluster node to the tree (see $\mathsf{CreateNewClusterNode}_{LD}$ below). The following rules define the action 'Create New Cluster' of the resource management layer and the respective rule it uses from the logical distribution (LD) layer:

CreateNewCluster$_{RM} \equiv$
  $newClusterNode(self) \leftarrow \mathsf{CreateNewClusterNode}_{LD}$

CreateNewClusterNode$_{LD} \equiv$
  let $nr = new(\textsc{Resource})$ **in**
    ConfigureResource$(nr, emptyCluster, noCapability)$
    CreateAndConfigureNode(nr)
    add $node(nr)$ **to** $childNodes(self)$
    $parent(node(nr)) := self$
    result $:= node(nr)$

The next step is to transfer the selected resources of the winning combination into the newly created cluster. Here, the resource manager layer directly uses the transfer capabilities service provided by the logical distribution layer (see $\mathsf{TransferCapabilities}_{LD}$ below). Based on the selected combination, a transfer map is created. Transfer of resources is done using a messaging protocol: for every pair of $(node, capability)$ in the given transfer map, a *transfer request* message is sent to *node* requesting to transfer its *capability* to the new cluster. Resources acknowledge a successful transfer of their capabilities by sending back a *transfer ack* message, which is collected in the next step.

The following rule defines the action 'Transfer Capabilities' of the logical distribution (LD) layer. Here we provide a simple implementation of the protocol. The node keeps a set of the capability transfer requests that it has sent so far, and then it removes those requests

for which an acknowledgement is received. The transfer is considered to be successful when all the requests are acknowledged before a time-out event occurs.

TransferCapabilities$_{LD}(task, target, transferMap) \equiv$
  if $requestsSent = undef$ **then**
    $outcome := undef$
    $timer \leftarrow$ GetNewTimer
    SendCapabilityTransferRequests$_{LD}(task, target, transferMap)$
  else
    if $|requestsSent| > 0$ **then**
      if $timedOut(timer)$ **then**
        $outcome := failed$
      else
        RemoveProcessedTransferRequests$(self, task, target, transferMap)$
    else
      $requestsSent := undef$
      $outcome := successful$
where
  $requestsSent \equiv capabilityTransferRequestsSent(self, task, target, transferMap)$
  $outcome \equiv transferResult(self, task, target, transferMap) = successful$
  $timer \equiv transferTimer(self, task, target, transferMap)$

SendCapabilityTransferRequests$_{LD}(task, target, transferMap) \equiv$
  seq
    $requestsSent := \{\}$
  next
    forall $(node, cap)$ **in** $transferMap$ **do**
      extend MESSAGE **with** m **do**
        $msgType(m) := transferRequest$
        $msgData(m,$ "capability" $) := cap$
        $msgData(m,$ "task" $) := task$
        $msgData(m,$ "target_node" $) := target$
        SendMessage$(node, m)$
        add $(node, cap)$ **to** $requestsSent$
where
  $requestsSent \equiv capabilityTransferRequestsSent(self, task, target, transferMap)$

8.1.5 The Executable Model

The DRCMA model is described in abstract functional and operational terms in form of an executable distributed abstract state machine specification, using the CoreASM modeling environment. This description of the underlying design concepts provides a concise blueprint for reasoning about key system attributes at an intuitive level of understanding, supporting requirements specification, design analysis and validation of system properties. A basic graphical user interface (see Figure 8.4) has also been developed in Java to provide a live view of the resource network and its command and control hierarchy during the simulation of scenarios. The specification utilizes the JASMine plugin (see Section 6.5) to interact with the graphical viewer.

As an example, Figure 8.4 shows the result of a simple search and rescue scenario using the network configuration of Figure 8.4(a) with the following resource capabilities and tasking requirements:

- Resource Capabilities
  - *Boat-1*: {{⟨mobility, high⟩, ⟨rescue, high⟩}}
  - *Aurora-1*: {{⟨mobility, high⟩, ⟨radar, medium⟩}}
  - *Helicopter*: {{⟨mobility, medium⟩, ⟨vision, medium⟩}}
  - *Boat-2*: {{⟨mobility, medium⟩, ⟨rescue, low⟩}}
  - *Aurora-2*: {{⟨mobility, high⟩, ⟨radar, high⟩}}
- Tasking Information
  - *SOS-task*: composed of *search-task* and *rescue-task*
  - Required capabilities
    - *search-task*: {{⟨mobility, medium⟩, ⟨vision, medium⟩}}
    - *rescue-task*: {{⟨mobility, medium⟩, ⟨rescue, low⟩}}

Building an abstract yet executable model of DRCMA in CoreASM enabled us to experiment with the model and validate design decisions at a fairly high level of abstraction. In subsequent steps, we are extending and further refining the DRCMA model into a comprehensive architecture [52] for adaptive distributed information fusion. The result will be a prototype for testing, experimental validation and machine-assisted verification of the key system attributes prior to actually building the system.

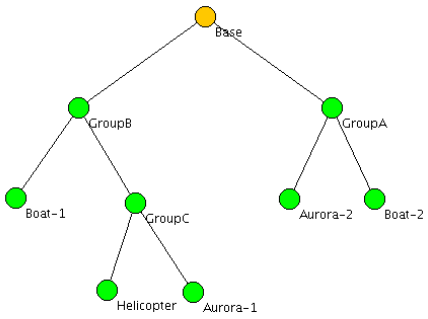

(a) Initial network configuration.

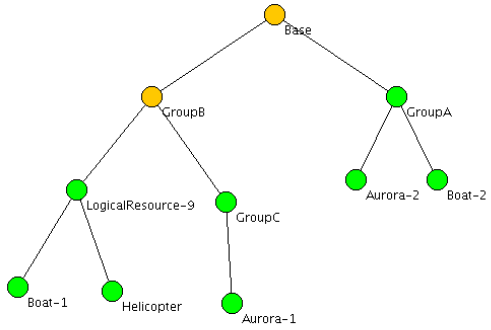

(b) A new cluster is created to perform the task.

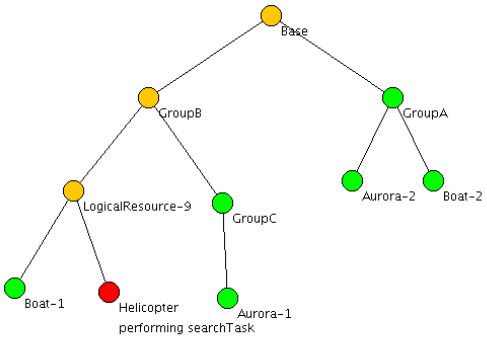

(c) Task is decomposed into two tasks and they are assigned to resources.

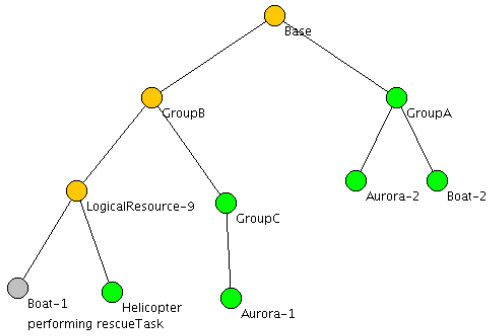

(d) Rescue task is assigned to Boat-1, but Boat-1 gets disconnected.

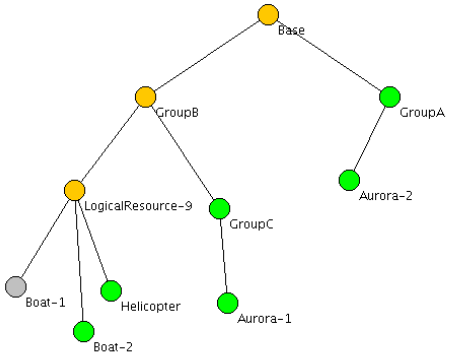

(e) Boat-2 is moved to the new cluster to perfrom the rescue task.

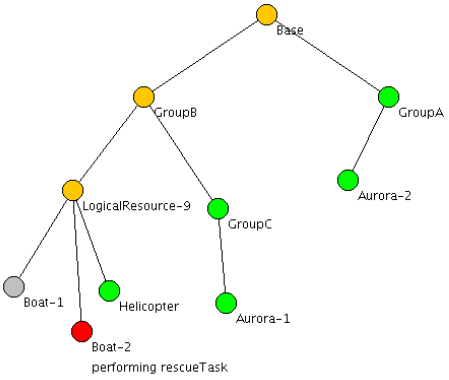

(f) Rescue task is assigned to Boat-2.

Figure 8.4: Search and Rescue Scenario

8.2 Decision Support for Situation Analysis

Situation Awareness is essential for conducting decision-making activities. It is the perception of elements in the environment, comprehension of their meaning, and projection of their status in the near future [43]. Agents develop an understanding of a situation based on a discrete perception and evaluation of events as they unfold over time and forecast their anticipated evolution in the future. Situation Analysis (SA) is defined as a process, the examination of a situation, its elements, and their relations, to provide and maintain a state of situation awareness for the decision maker [28].

The rationale for establishing a formal semantic foundation for the design of situation analysis and decision support systems is discussed in detail in [27]. Inspired by recent work at Defence R&D Canada at Valcartier that proposes the use of Interpreted Systems (see Section 2.4) for Situation Analysis [94, 95], a systematic approach combining ASMs and Interpreted Systems seems appealing, as each of the two semantic modeling frameworks has its particular focus and strength, complementing each other in several respects. They both share common abstraction principles for describing distributed system behavior based on an abstract operational view of multiagent systems. Additionally, pragmatic considerations regarding practical needs for system design and development are relevant to support the systematic refinement of abstract specifications into executable models serving as a basis for rapid prototyping and experimental validation of decision support systems.

In the following sections (previously published in [50]), we illustrate the similarities between the two frameworks using a simple surveillance scenario originally presented in [95]. In order to put the abstract model into practice and to realize what is practically feasible, we used the CoreASM modeling suite to produce an executable model of the scenario through refinement of the abstract rules and functions of the model. Such a refinement, with the goal of producing an executable model, is interesting in two aspects: *a)* it helps in finding ambiguities, missing pieces and loose-ends of the model and forces the system analyst/modeler to think clearly about the main concepts and their definitions, and *b)* it supports experimental validation through execution (simulation).

Our work demonstrates how one can benefit from using ASMs, and in particular CoreASM, to model a multiagent system while still being able to apply and extend the Interpreted Systems approach of [95] for situation analysis. For example, by combining Interpreted Systems and multiagent ASMs, situation analysis queries can not only be analyzed using the

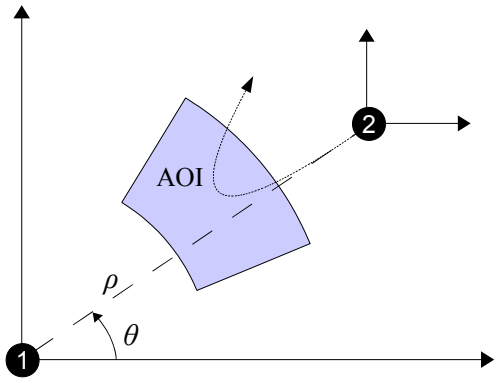

Figure 8.5: Surveillance Scenario from [95] [6]

proposed methods in [95], but they can also be examined either by explicitly encoding the queries as computable functions in the model and running the executable model, or by applying the available model checking techniques for ASM [93, 66].

8.2.1 The Abstract Model

The scenario involves two agents, $agent_1$ and $agent_2$, in a 2D environment (see Figure 8.5). Both agents are able to move in 8 possible directions and they can both sense other agent's range ρ and bearing θ with some error. The local state of each agent is composed of successive observations about the other agent's position. The state of the environment contains the accurate target positions. Using the ASM notation, we have:

universe AGENT $= \{agent_1, agent_2, env\}$
universe MOVEMENT $= \{\mathrm{N}, \mathrm{NW}, \mathrm{W}, \mathrm{WS}, \mathrm{S}, \mathrm{SE}, \mathrm{E}, \mathrm{EN}\}$

$\rho,\ \theta : \text{AGENT} \mapsto \mathbb{R}$	$\hat{\rho}_{1h},\ \hat{\theta}_{1h} : \ \mapsto \text{SEQUENCE}(\mathbb{R})$
	$\hat{\rho}_{2h},\ \hat{\theta}_{2h} : \ \mapsto \text{SEQUENCE}(\mathbb{R})$
$\hat{\rho}_1,\ \hat{\theta}_1 : \ \mapsto \mathbb{R}$	$error_\rho : \text{AGENT} \mapsto \mathbb{R}$
$\hat{\rho}_2,\ \hat{\theta}_2 : \ \mapsto \mathbb{R}$	$error_\theta : \text{AGENT} \mapsto \mathbb{R}$

where $\rho(a)$ and $\theta(a)$ are the real range and bearing of agent a, $\hat{\rho}_i$ and $\hat{\theta}_i$ are the observed range and bearing[7] and $\hat{\rho}_{ih}$ and $\hat{\theta}_{ih}$ are the history of observations of range and bearing of

[6] Courtesy of A.-L. Jousselme and P. Maupin

[7] A better approach would be to define $\hat{\rho}$ and $\hat{\theta}$ as functions over agents, but to match the original syntax of the scenario, we define them as individual indexed functions.

agent i by the other agent, and $error_\rho(a)$ and $error_\theta(a)$ are the observation errors of agent a.

To make the model more general and scalable, real ranges and bearings are defined as functions over agents. The action sets of the two agents, ACT_1 and ACT_2 are equivalent to the Movement universe defined above.

According to the scenario, $agent_1$ is stationary. Its purpose is to observe the position of $agent_2$ and to send a message to a designated agent (outside of the system) if $agent_2$ enters an area of interest (AOI). On the other hand, $agent_2$ is actively moving toward $agent_1$ until it finds out that it is "too close" to $agent_2$, in which case it makes a U-turn.[8] So, we define the programs of $agent_1$ and $agent_2$ as follows:

Agent1Program $\equiv$
  RecordObservation($agent_2$)
  if $isInAOI(agent_2)$ **then**
    SendMessage("Agent 2 is in AOI.")

Agent2Program $\equiv$
  RecordObservation($agent_1$)
  if $dir(self) = toward$ **then**
    MoveToward($agent_1$)
  else
    MoveAway($agent_1$)
  if $tooClose(agent_1)$ **then**
    $dir(self) := away$

At this level, we abstract away from certain details. For example, since the scenario does not elaborate on how and to whom $agent_1$ sends its messages, we leave the SendMessage rule abstract. Also, since the area of interest and the exact measures for "closeness" of $agent_1$ and $agent_2$ are not defined, the functions $isInAOI$ and $tooClose$ are left abstract as well. The rule RecordObservation(a) is a placeholder for the actual task of maintaining the observation history.

An interesting observation is that the local state of $agent_2$, as described by the scenario, seems to be missing an important piece: the movement direction of $agent_2$, toward or away

[8] The description of the scenario is vague on whether $agent_2$ will ever try to move back toward $agent_1$ or not. We assume that it will not try to come close again.

from $agent_1$. At some point in time, $agent_2$ observes that it is too close to $agent_1$ and so makes a U-turn. The direction of the $agent_2$ is a dynamic attribute of the agent and needs to be captured in the local state of the agent. Here, we model the direction of $agent_2$ by a function *dir*, which is initially set to *toward* and will be changed to *away* the first time $agent_2$ observes that it is too close to $agent_1$.[9]

In this scenario, the environment models the uncertainty of error values. The action of the environment is to set the error values for observed ranges and bearings. Hence, the action set of the environment ACT_e is a set of tuples of the form $(e^1_\rho, e^1_\theta, e^2_\rho, e^2_\theta)$ in which $e^i_\rho \in E^i_\rho$ and $e^i_\theta \in E^i_\theta$ are range and bearing observation errors for agent i, and E^i_ρ and E^i_θ are the corresponding error ranges. In our ASM model, the environment program models this behavior by non-deterministically choosing values from E^i_ρ and E^i_θ and updating $error_\rho$ and $error_\theta$ for agents $agent_1$ and $agent_2$.

EnvironmentProgram ≡
  forall $a \in \{agent_1, agent_2\}$ **do**
    choose $e_\rho \in E_\rho(a),\ e_\theta \in E_\theta(a)$ **do**
      $error_\rho(a) := e_\rho$
      $error_\theta(a) := e_\theta$

At this point, we have an abstract operational model of the scenario in form of a multi-agent ASM.

8.2.2 Situation Awareness

Following the approach of [95], let Φ be the basic set of our propositions, $\phi_\rho \in \Phi$ be "*the range of agent 2 crosses AOI*", and $\phi_\theta \in \Phi$ be "*the bearing of agent 2 crosses AOI*". The formula $\phi_{\text{AOI}} = \phi_\rho \wedge \phi_\theta$ will then be "*agent 2 is in AOI*". In the program of $agent_1$, the function $isInAOI(agent_2)$ represents the awareness of $agent_1$ about ϕ_{AOI}, and it can be evaluated using the range and bearing of $agent_2$ as observed by $agent_1$:

$$isInAOI(agent_2) = \hat{\rho}_2 \in AOI_\rho \wedge \hat{\theta}_2 \in AOI_\theta \tag{8.1}$$

[9] One can also argue that the knowledge about the close encounter of the agents is implicitly encoded in the observation history of $agent_2$ and the direction can be dynamically calculated based on the observation history.

It is important to note that since there is a possible error in the observation of $agent_1$, $isInAOI(agent_2)$ being *true* does not necessarily mean that ϕ_{AOI} holds as well; i.e., $agent_2$ could actually be outside of the area of interest. The value of $isInAOI(agent_2)$ simply reflects the state of awareness of $agent_1$ about the position of $agent_2$, which may differ from the reality. This is captured in (8.1) by using the observed range and bearing, $\hat{\rho}$ and $\hat{\theta}$, rather than the real values ρ and θ.

The same holds for the awareness of $agent_2$ about its distance from $agent_1$. Let $\phi_c \in \Phi$ be "*the range of agent 1 is too close*". The function $tooClose(agent_1)$ represents the state of awareness of $agent_2$ about the truth of ϕ_c and can be derived from the observed range of $agent_1$:

$$tooClose(agent_1) = \hat{\rho}_1 < threshold_\rho^2 \tag{8.2}$$

where $threshold_\rho^2$ is the minimum distance that $agent_2$ is willing to have with $agent_1$. Again, $tooClose(agent_1)$ being *true* does not necessarily mean that ϕ_c holds.

We can further extend the model and introduce non-trivial computable formulas such as $\phi_m =$ "*agent 2 is coming toward agent 1*". In general, derived functions can be defined to model the awareness of agents about the truth values of such formulas. For example, the following function can represent the awareness of $agent_1$ about ϕ_m:

$$approaching_2 = \hat{\rho}_{2h}(last) < \hat{\rho}_{2h}(last - 1) \tag{8.3}$$

8.2.3 Situation Analysis

Once we have a proper model of the scenario, various types of queries can be used for situation analysis. Using the Interpreted Systems framework, Jousselme and Maupin suggest the following three types of queries for situation analysis [95]:

1. Queries about truth, such as "Does ϕ_{AOI} hold in a given state s?";
2. Queries about knowledge, such as "Does $agent_2$ know that ϕ_c holds in a given state s?"; and
3. Queries about time, such as "Does ϕ_{AOI} eventually hold in a run r of the system?"

By combining Interpreted Systems and multiagent ASMs, these queries can not only be analyzed using the proposed methods in [95], but they can also be examined either by explicitly encoding the queries as derived functions (such as *isInAOI* and *tooClose*) and

running the executable model, or by applying the available model checking techniques for ASM [93, 66]. Thus, the approach of integrating ASM with Interpreted Systems is consistent with the proposed approach of [95]. Note that the idea of using model checking for situation analysis has first been proposed in [95].

8.2.4 Executable Model

In order to put the model into practice and to realize what is practically feasible, one needs to experiment with the model; that is, the model has to be machine executable.

In this section, we produce an executable model of the scenario through refinement of the abstract rules and functions of the model presented in Section 8.2.1. We then use the CoreASM execution engine to run the model. Such a refinement, with the goal of producing an executable model, is interesting in two aspects: *a)* it helps finding ambiguities, missing pieces and loose-ends of the model and forces the system analyst/modeler to think clearly about the main concepts and their definitions, and *b)* it supports experimental validation through execution (simulation).

Real Positions and Rules of Movement

The refinement of the routines MoveToward and MoveAway requires a proper encoding of the positions of the agents. Although it is not precisely stated in the scenario, the real positions of the agents appear to be relative to the positions of their observers. Regardless of the encoding one chooses, the observed positions and the rules of movement both depend on the real position values and must be defined consistently. In our refined model, we keep the actual position of every agent in an (x, y) coordinate which makes it easier to define the movement routines. Real relative bearing and range values are simply calculated based on the actual positions of agents.

Observed Values

An important piece that was left abstract in our model, and was also missing in the original scenario, is the definition of the observed range and bearing functions $\hat{\rho}_i$ and $\hat{\theta}_i$. These functions play an important role in the model and their formal definition is important for proper situation analysis. While the original model does not state how the observed values

are produced, it is clear that the observed position values of agent i are functions of both the real position of agent i and the corresponding observation error. Hence, the following equations are reasonable candidates:

$$\hat{\rho}_i = \rho(agent_i) + e^i_\rho \ , \quad \hat{\theta}_i = \theta(agent_i) + e^i_\theta$$

So, the observation functions can be defined as derived functions over the actual positions of the agents. Another approach, though less intuitive, would be to have the environment explicitly update the observed values in every step of the system.

Recording Observation

Since we define the observation values as derived functions over the real positions of the agents, to record the observation history we simply add the current values of $\hat{\rho}_i$ and $\hat{\theta}_i$ to the observation histories of each agent:

RecordObservation$(agent_i) \equiv$
  add $\hat{\rho}_i$ **to** $\hat{\rho}_{ih}$
  add $\hat{\theta}_i$ **to** $\hat{\theta}_{ih}$

At this point, we have a machine executable model of the scenario (see Appendix B.2) and we can use the CoreASM execution engine to run a simulation of the model. We start with an initial state in which $agent_1$ is located at (0, 0) and $agent_2$ is located at (15, 10). To monitor the position of $agent_2$, we extend the program of the environment to print the current positions of the agents in every step. This is a sample output of the simulation:

```
agent1:(0, 0) - agent2:(15, 10)
agent1:(0, 0) - agent2:(14, 9)
agent1:(0, 0) - agent2:(13, 8)
agent1:(0, 0) - agent2:(12, 7)
agent1:(0, 0) - agent2:(11, 6)
agent1:(0, 0) - agent2:(10, 5)
Abstract Call: SendMessage(Agent 2 is in AOI.)
agent1:(0, 0) - agent2:(11, 5)
agent1:(0, 0) - agent2:(10, 6)
Abstract Call: SendMessage(Agent 2 is in AOI.)
```

```
agent1:(0, 0) - agent2:(9, 7)
Abstract Call: SendMessage(Agent 2 is in AOI.)
agent1:(0, 0) - agent2:(8, 8)
agent1:(0, 0) - agent2:(7, 9)
agent1:(0, 0) - agent2:(6, 10)
agent1:(0, 0) - agent2:(5, 11)
agent1:(0, 0) - agent2:(4, 12)
agent1:(0, 0) - agent2:(4, 13)
agent1:(0, 0) - agent2:(4, 14)
```

The simulation starts with $agent_2$ moving toward $agent_1$. When $agent_1$ observes that $agent_2$ is in the area of interest, it sends a message using the abstract routine SendMessage (which is intentionally left abstract). After some time, $agent_2$ realizes that it is too close to $agent_1$ and makes a U-turn, moving away from $agent_1$.

8.3 The Mastermind Project

The Mastermind project [32, 31, 51] is a pioneering interdisciplinary project in computational criminology that employs formal modeling and simulation as tools to investigate offender behavior in urban environments.[10] It is jointly managed by the Software Technology Lab and the Institute for Canadian Urban Research Studies (ICURS)[11] at Simon Fraser University, aiming at developing computational models of criminal activity patterns, with a special focus on spatio-temporal characteristics of crime, potentially involving multiple offenders and multiple targets. The Mastermind project utilizes the ASM method and the CoreASM tool suite to address the specific requirements of developing computational models and analysis tools for the study of crime in a collaborative research environment.

Crime is composed of four main elements: the law, the offender, the target and the location [29]. The Mastermind project constructs a multi-dimensional model of crime in order to study the interaction of these elements. The focus is on the concepts of environmental criminology, which argues that in spite of their complexity, criminal events can be understood in the context of people's movements in the course of everyday routines [29]. Through movement within a given environment, possible offenders, characterized as *agents*, develop

[10] The Mastermind project is not a contribution of this thesis. It is addressed here only as an example of the application of CoreASM in interdisciplinary projects.

[11] http://www.sfu.ca/icurs

mental maps of the places they know (*awareness space*) and the places they regularly visit (*activity space*). At its core, Mastermind captures the essence of the Crime Pattern theory that states: crime occurs when a motivated individual encounters a suitable target [29]. Figure 7.5(b) captures this behavior in terms of a Control State ASM.

At the heart of the Mastermind project is a robust ASM ground model [26] developed through several iterations required for checking the validity of the model with respect to the understanding of domain experts. The process of establishing the key properties, determining the right level of abstraction, and ensuring the validity of the model was greatly facilitated using the simple graphical notation provided by CSDe and the ability to run experiments on abstract models in early stages of design using the CoreASM engine.

The ground model has been further refined into more concrete models with specific details systematically added. The simulation model of Mastermind implemented in Java is an example of such refinements. The Java version provides a graphical user interface and a simulation environment based on real-world Geographical Information System (GIS) data and captures the navigation behavior of offenders with a high degree of detail and complexity. The CoreASM executable ground model has also been further refined to run more controlled experiments in CoreASM, which allows for a structured analysis of theories in a hypothetical world. A special plugin for CoreASM is developed offering a custom visualization of the simulation (see Figure 8.6). These simple and comprehensible models provide domain experts with full control over the variables under study and their interdependence. Both the Java and CoreASM model provide visualization features which are a priority for criminology publications. Figure 8.6 shows a snapshot of the Mastermind plugin for CoreASM. The visualization shows the movement of agents between activity nodes, the formation of their activity spaces and the effects on crime hotspots.

According to [30, 51], CoreASM has played an important role in facing the challenges of two major phases of the Mastermind project, namely *formalization* and *validation*. In an interdisciplinary research project, the communication problem is intensified, imposing serious challenges in ensuring a correct transformation from domain knowledge to computational artifacts. The differences between academic disciplines in terms of approach and underlying assumptions, and the fact that real-life events, such as crime events, are not usually thought of in a discrete, mathematical manner, further complicate the communication issues. To this end, diagrams created by CSDe greatly facilitate an interactive design process where domain experts are able to directly check and correct a design.

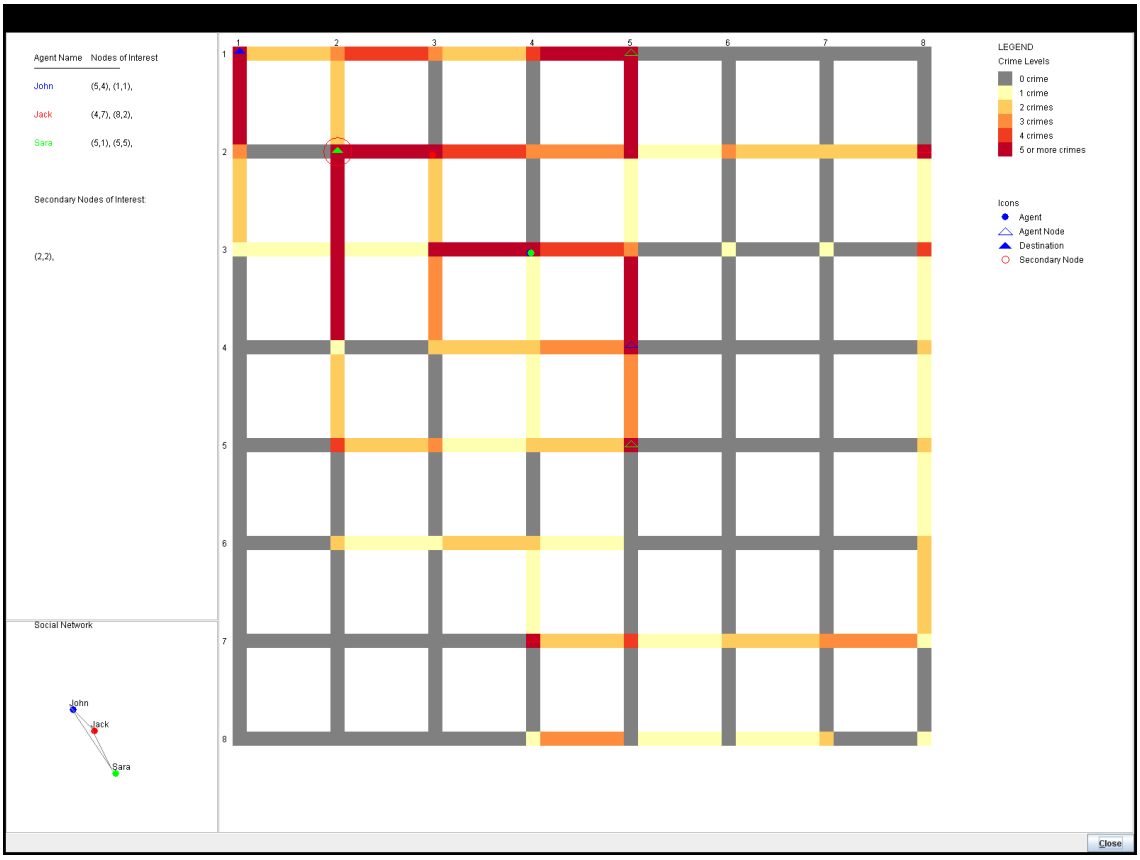

Figure 8.6: The Mastermind Plugin for CoreASM

As a final remark, it is important to compare the utility of the full-fledged Mastermind simulation model in Java with the simpler, more abstract CoreASM model. The complexity of the Java version and the fact that it is considered as a black-box by domain experts introduces limitations on its academic use. On the other hand, the CoreASM program code is easier for non-programmers to read, and is well-suited for designing *controlled* experiments. Taking advantage of the highly flexible plugin architecture offered by CoreASM, the team were able to rapidly develop the Mastermind Plugin to address the specific needs of criminologists, especially with respect to visualizing the results. In other words, the Mastermind Plugin encapsulates the mathematical structure of the ASM model in a comprehensible and familiar format for domain experts. This greatly facilitates communication with domain experts and analysis of the results for validation purposes.

Chapter 9

Conclusions and Perspectives

This work presented the design and development of the CoreASM modeling framework and tool environment for high-level design and analysis of abstract state machine models. The CoreASM engine forms the kernel of a novel environment for model-based engineering of abstract operational requirements and design specifications at the early phases of the software design and development process. Focusing on *freedom of experimentation* and *design exploration*, CoreASM offers a flexible modeling environment that facilitates writing of easily modifiable, concise and understandable formal specifications by minimizing the need for encoding of domain concepts into the constructs of the language.

In order to minimize the cost of such encoding, the CoreASM language and tool architecture are both designed to be easily extensible so that they can be customized for specific application contexts, thus realizing *domain-specific* ASM dialects. The ASM literature contains many examples of using such ASM dialects: many published specifications of large systems have introduced background elements or non-standard rule forms that were well suited to express the intended behavior at an appropriate level of abstraction in the given domain. By similarly allowing the customization of the CoreASM language, we provide the benefits of executable specifications without loosing the expressiveness of a domain-specific language, and avoid the introduction of a further encoding level between the conceptual specification and its executable version.

The design of the CoreASM engine is formally specified in ASMs. The entire lifecycle of the CoreASM engine is defined as an extensible control-state ASM and the CoreASM language is formally defined through the specification of an interpreter (in the form of an abstract state machine) that ensures the executability of the language and provides its

formal semantics.

CoreASM has been recognized by the ASM community and has been used by various research groups in Europe, Asia, and North America [91, 3, 6, 85, 97, 40].[1] Based on solid experience gained through the practical use of CoreASM in a number of diverse application domains (see Chapter 8), we claim that CoreASM serves practical needs of high-level modeling and rapid prototyping of complex distributed systems and will be an asset for industrial engineering of complex software systems by making software specifications and designs more robust and reliable. Prior to actually building a system, CoreASM specifications facilitate development of concise blueprints for intuitive reasoning about key system attributes, supporting requirements specification, design analysis, validation and (where appropriate) formal verification of system properties.

9.1 Significance of the Contribution

Among all the existing ASM tool environments, CoreASM stands out as being the closest to the spirit of abstract state machines [25]. Here, we summarize the most significant features that distinguish it from other ASM tools.

A Rich ASM Language and Framework

CoreASM offers a rich ASM language with a syntax that closely follows the pseudo-code style of ASMs and a formally defined semantics that is faithful to the original ASM semantics as defined in [25]. CoreASM is the first ASM tool environment that directly supports distributed ASM computation models with custom scheduling policies. Its language supports classes of basic, distributed, and Turbo ASMs, making it the most comprehensive ASM language available.

Encouraging Rapid Prototyping

The CoreASM language is an untyped language with a minimal yet human-readable syntax that facilitates writing abstract and untyped models which can be refined into more concrete

[1] To name a few, CoreASM has been applied in a number of research projects at the Computer Science Department of the University of Pisa in Italy, the Embedded Software Laboratory at the RWTH Aachen University in Germany, the Open Systems Development Group at the University of Agder in Norway, and the Department of Computer Science and Engineering at the Anna University in India.

versions as needed. Thus, it encourages rapid prototyping of abstract machine models for testing and design space exploration, and facilitates agile software development [64]. An independent study performed by Jensen et al. [85], comparing the abstraction level of specifications written in CoreASM and AsmL[2], shows that the CoreASM language can be used to specify algorithms in a higher level of abstraction compared to AsmL. In their example of a data clustering algorithm, the CoreASM description of the algorithm is 82 lines, almost half the size of the 155 lines of AsmL description of the same algorithm (see [85, Fig. 4]). The authors conclude that compared to AsmL, CoreASM is more suited for the early stages of software engineering.

Extensible Language and Architecture

The most significant feature of CoreASM is the extensibility of its language and modeling environment. To reduce the cost of writing specifications, one has to minimize the need for encoding in mapping the problem space to a formal model. This approach usually leads to the design of domain-specific languages. The CoreASM extensibility framework provides utmost flexibility for extending its language definition and execution engine in order to tailor it to the particular needs of virtually any conceivable application context. This allows CoreASM to be used very much in the same way ASMs were meant to be used.

An Open Framework

CoreASM is one of the few ASM tools that is implemented as an open framework. Developed in Java—a platform independent, open source programming language—and under an open source license, CoreASM can be modified, extended and improved as needed by its user community. The CoreASM engine comes with a simple yet comprehensive API that offers full access to the states of simulated machines and complete control over the execution of CoreASM specifications, and as such facilitates the integration of CoreASM as an ASM simulator component into other applications.

[2]The executable ASM language developed by the Foundation of Software Engineering group at Microsoft [101]

9.2 Future Work

The CoreASM project is in continuous development. Currently, the execution engine can execute standard ASM specifications; various plugins offer common backgrounds such as numbers, sets, strings, and lists, and more specialized plugins offer sophisticated features such as the JASMine plugin for interfacing ASM specifications with Java class libraries (see Section 6.5).

However, there are a number of open issues that have not been yet sufficiently addressed by the CoreASM project. In this section we review some of these issues and discuss them as possible subjects of future work.

Debugging Features

Traditional debugging models of programming (e.g. step by step execution of instructions) do not suit ASMs. There is no such concept as the "current" instruction, nor an explicit notion of "stepping" over instructions. However, similar notions can be applied to computation steps of ASMs instead.

For example, a debugging user interface can offer, after every step of simulation, the option of browsing the ASM program as a tree of rule constructs annotated with the most recently generated update multisets produced by the rules. Such a feature would allow users to investigate the changes (updates) produced by different parts of ASM programs at desired levels of detail.

The CoreASM engine provides the necessary services (such as step-by-step execution of the engine, full access to the simulated state, and the possibility of applications to intervene in the execution process of the engine) supporting the implementation of various debugging features by a CoreASM user interface. Non-trivial debugging features, however, are not yet implemented in any of the currently available CoreASM user interfaces.

Type System

The CoreASM language is designed as a primarily typeless language to encourage rapid prototyping of abstract specifications. Although dynamic types are attached to every CoreASM value (element) and various primitive and complex data types are provided by plugins (see sections 6.2 and 6.3), there is no concept of static typing or a type system defined in the

kernel of CoreASM. State locations (a more generalized notion of programming variables) are essentially typeless and there is no type-checking offerred by the CoreASM kernel.

The Signature plugin (see Section 6.4.1) extends the CoreASM language and the engine by offering a means to define type signatures for state locations. It also provides runtime type checking on function calls and on updates to locations for which a signature is defined. However, much more can be done in this domain. For example, collection plugins could be improved to offer parameterized type constructors and the Signature plugin could be extended to offer static type analysis of fully-typed specifications, a practical requirement for model checking of CoreASM specifications.

Literate Specifications

Following the idea of *literate specifications* [86] (an extension of Knuth's *literate programming technique* [89]), it would be beneficial to integrate facilities for writing CoreASM specifications into various document preparation systems such as OpenOffice Writer[3] or the LaTeX typesetting system[4]. Such an integration would facilitate the development of compound system documents, consisting of executable specifications and system documentations, that not only provide formal specification of systems, but also offer design rationale and necessary explanation on how such systems work.

The current implementation of CoreASM can import specifications from OpenOffice Writer documents and the Carma user interface (see Section 7.3) can load and execute OpenOffice Writer documents containing CoreASM fragments. The CoreASM engine could be extended to also support import and export of specifications to and from LaTeX documents.[5]

Integrated Development Environment

The CoreASM IDE, a combination of the CoreASM Eclipse plugin and the CSD editor (see Section 7.3), is still in early development. We envision further improvements providing debugging features (discussed above) and enhanced coding assistance features, such as easy

[3] http://www.openoffice.org/

[4] http://www.latex-project.org/

[5] A basic CoreASM-to-LaTeX export feature has already been implemented in Carma which has been used to produced the color-annotated specification of Appendix B.1.

navigation between different layers of abstraction and refinements, which would be of real value in building complex models.

Verification and Model Checking

A proper formal specification facilitates establishing the validity of the initial formalization step, which itself is a prerequisite for any meaningful approach to formal verification. However, the only machine-assisted verification supported by the current implementation of CoreASM is in the form of rudimentary model checking (see [93] and Section 7.3.2). More sophisticated interfaces to existing model checking tools are needed to fully exploit the potential they provide.

Automatic Code and Test Case Generation

There is currently no support for automatic code generation from CoreASM models. The CoreASM engine is reasonably fast and efficient for interactive modeling and experimental validation; nonetheless, there is room for improving performance by generating Java or C++ code from CoreASM specifications. Automatic test case generation for conformance testing, comparable to AsmL Spec Explorer [122], is a work in progress independent of our work.

Part IV

Appendices

Appendix A

Supplementary Definitions

A.1 Abstract Storage

- PushState puts the current state in the stack. We assume that $stack_{state}$ is empty in the initial state.

  PushState $\equiv$
  Push$(stack_{state}, state)$

- PopState retrieves the state from the top of the stack, thus discarding the current state.

  PopState $\equiv$
  $state := top(stack_{state})$
  Pop$(stack_{state})$

- Apply(u) applies the updates in the update set u to the current state.

  Apply(u) $\equiv$
  forall $(l, v) \in u$ **do**
  SetValue(l, v)

- ClearState resets *state* to an empty state.

  ClearState $\equiv$
  let $s = new(\textsc{State})$ **in**
  $state := s$

- *newElement* : Element
  returns a new element; i.e., imports a new element into the state and returns the imported element. This function is defined as follows:

$$newElement \equiv new(\text{Element})$$

- *inconsistentUpdates* : Set(Update) $\mapsto$ Set(Update)
  returns the set of inconsistent updates (according to [25, Def. 2.4.5]) in the given update set. We assume that the update set consists of regular updates only (i.e. actions are *updateAction*).

$$inconsistentUpdates(uset) \equiv \{(l, v, a) \in uset \mid \exists (l', v', a') \in uset,\ l = l' \wedge v \neq v'\}$$

- *isConsistent* : Set(Update) $\mapsto$ Boolean
  returns *true* if the update set is consistent according to [25, Def. 2.4.5]. We assume that the update set consists of regular updates only (i.e. actions are *updateAction*).

$$isConsistent(uset) \equiv |inconsistentUpdates(uset)| > 0$$

- *isUniverseName* : Name $\mapsto$ Boolean

$$isUniverseName(name) \equiv universes(state, name) \neq undef$$

- *isFunctionName* : Name $\mapsto$ Boolean

$$isFunctionName(name) \equiv functions(state, name) \neq undef$$

- *isRuleName* : Name $\mapsto$ Boolean

$$isRuleName(name) \equiv rules(state, name) \neq undef$$

A.2 Interpreter

- ClearTree(t) clears the given tree from any assigned value, location, or updates.

ClearTree(α) $\equiv$
  if $\alpha \neq$ *undef* **then**
    value(α) := *undef*
    update(α) := *undef*
    loc(α) := *undef*
    ClearTree(*first*(α))
    ClearTree(*next*(α))

- CopyTree(t, *setNext*) creates a copy of the given tree, without copying assigned values, locations, or updates. If *setNext* is true, it also copies the next sibling of the given root node.

CopyTree(α, *setNext*) $\equiv$
  if $\alpha \neq$ *undef* **then**
    let n = *new*(NODE) **in**
      class(n) := *class*(α)
      pattern(n) := *pattern*(α)
      token(n) := *token*(α)
      grammarRule(n) := *grammarRule*(α)
      plugin(n) := *plugin*(α)
      first(n) := CopyTree(*first*(α), *true*)
      if *setNext* **then**
        next(n) := CopyTree(*next*(α), *true*)
      result := n
  else
    result := *undef*

- CopyTreeSub($\alpha, \langle x_1, \ldots, x_n \rangle, \langle \lambda_1, \ldots, \lambda_n \rangle$) returns a copy of the given parse tree α, where every instance of a parameter node x_i is substituted by a copy of the corresponding argument λ_i. We assume that the elements in the formal parameters list (x_i's) are all distinct. Also, formal parameters substitution is applied only to occurrences of formal parameters in the original tree passed as argument, and *not* also on the actual parameters themselves.

CopyTreeSub$(\alpha, \langle x_1, \ldots, x_n \rangle, \langle \lambda_1, \ldots, \lambda_n \rangle) \equiv$
  if $\alpha \neq$ *undef* **then**
    if *class*$(\alpha) =$ Id $\wedge \exists i$ s.t. *token*$(\alpha) = x_i$ **then**
      result $\leftarrow$ CopyTree$(\lambda_i, false)$
    else
      let $n =$ *new*(NODE) **in**
        first$(n) \leftarrow$ CopyTreeSub(*first*$(\alpha), \langle x_1, \ldots, x_n \rangle, \langle \lambda_1, \ldots, \lambda_n \rangle)$
        next$(n) \leftarrow$ CopyTreeSub(*next*$(\alpha), \langle x_1, \ldots, x_n \rangle, \langle \lambda_1, \ldots, \lambda_n \rangle)$
        class$(n) :=$ *class*(α)
        pattern$(n) :=$ *pattern*(α)
        token$(n) :=$ *token*(α)
        grammarRule$(n) :=$ *grammarRule*(α)
        plugin$(n) :=$ *plugin*(α)
        result $:= n$
  else
    result := *undef*

- HandleUndefinedIdentifier(*pos*, *x*, *args*) asks all the plugins registered to handle undefined identifiers to evaluate the node with the undefined identifier (*pos*). It is considered an error if more than one plugin evaluates the undefined identifier with different results. If none of the plugins could evaluate the node, KernelHandleUndefIdentifier will be called to create a new function element with a default value of undef_e for the given arguments.

HandleUndefinedIdentifier$(pos, x, args) \equiv$
  local *results* [*results* := {}] **in**
    seq
      foreach p **in** *loadedPlugins* **do**
        seqblock
          ClearTree(*pos*)
          PluginHandleUndefIndentifier(p, *pos*, *x*, *args*)
          if *evaluated*(*pos*) **then**
            add $\langle p, loc(pos), updates(pos), value(pos) \rangle$ **to** *results*
        endseqblock

next
 if $|results| = 0$ **then**
  KernelHandleUndefIdentifier($pos, x, args$)
 else
  choose $\langle p, l, u, v\rangle$ **in** *results* **with** $\exists \langle p', l', u', v'\rangle \in results, \langle l, v, u\rangle \neq \langle l', v', u'\rangle$ **do**
   Error('There is an ambiguity in resolving the identifier.')
  ifnone
   $[\![pos]\!] := (l, u, v)$

A.3 Scheduler

- *updateInstructions* : Multiset(Update)
  is the multiset of accumulated update instructions in the current computation step.

- *updateSet* : Set(Update)
  is the set of (aggregated) updates in last computation step.

- *selectedAgentsSet* : Set(Element)
  is the set of selected agents contributing to the computation of the current step.

- *initAgent* : Element
  is the initial agent the engine creates to run the *init* rule.

- *chosenAgent* : Element
  is the currently running (or to be running) agent.

- *chosenProgram* : Rule
  is the rule element that represents the program of the chosen agent. The value of this function is set by the Abstract Storage.

- *morePossibleSetsExist* : Boolean
  holds true if there are more possible combinations of agents that can contribute to the current computation step.

- *isSingleAgentInconsistent* : Boolean
  holds true if the last inconsistent set of updates is produced by a single agent.

$$isSingleAgentInconsistent \equiv$$
$$\exists a \in \textsc{Element}, \exists l \in \textsc{Location}, \forall u_1, u_2 \in updateSet,$$
$$uiLoc(u_1) = uiLoc(u_2) \wedge uiAgents(u_1) = uiAgents(u_2) = \{a\}$$

- LoadSchedulingPolicy, based on the set of loaded plugins, loads a scheduling policy for scheduling of agents in every computation step.

```
LoadSchedulingPolicy ≡
  let policies = {pluginSchedulingPolicy(p) | p ∈ specPlugins ∧ isPolicyPlugin(p)}\{undef} in
    if |policies| = 0 then
      schedulingPolicy := undef
    else
      if |policies| = 1 then
        choose policy ∈ policies do
          schedulingPolicy := policy
          schedulingGroup := newSchedulingGroup(policy)
      else
        Error('Conflicting scheduling policies.')
```

A.4 Control API

The following functions and rules define the interface of the engine to its environment.

- *specification* : Spec
  is the current CoreASM specification loaded by the engine.

- *pluginCatalog* : Set(Plugin)
  is the set of all the plugins available to the engine.

- *loadedPlugins* : Set(Plugin)
  is the set of loaded plugins by the engine.

- *grammarRules* : Set(GrammarRule)
  is the set of all the grammar rules provided by the kernel and loaded plugins.

- *isStateInitialized* : Boolean
  holds true if the simulation state is initialized.

- *stepCount* : Number
  is the simulation step counter.

- *state* : State
  holds the current simulation state.

- *agentSet* : Set(Element)
  is the set of all the available agents in the current state retrieved from the Abstract Storage at the beginning of every computation step.

- *engineProperties* : Name $\mapsto$ Name
  holds all the defined engine properties and their values. The behavior of the engine (and its plugins) can be customized by these properties.

- *engineMode* : EngineMode
  returns the current execution mode of the engine.

- *isEngineBusy* : Boolean

$$isEngineBusy \equiv engineMode \notin \{Idle, Error\}$$

- UpdateState(*updates*), if $\neg isEngineBusy$, updates the current state by applying the given set of updates.

- Step puts a *step* command in the command queue of the engine.

A.5 Plugins

A.5.1 Choose Rule Plugin

Choose Rule

$(\!\|\textbf{choose}\ ^{\alpha}x\ \textbf{in}\ ^{\beta}\boxed{e}\ \ \textbf{do}^{\gamma}\boxed{r}\ \|\!)$	$\rightarrow$	$pos := \beta$
$(\!\|\textbf{choose}\ ^{\alpha}x\ \textbf{in}\ ^{\beta}v\ \textbf{do}^{\gamma}\boxed{r}\ \|\!)$	$\rightarrow$	**if** $enumerable(v)$ **then**
		let $s = enumerate(v)$ **in**
		if $\|s\| > 0$ **then**
		choose $t \in s$ **do**
		$\mathsf{AddEnv}(x, t)$
		$pos := \gamma$
		else
		$[\![pos]\!] := (undef, \{\!\}, undef)$
		else
		Error('Cannot choose from a non-enumerable element.')
$(\!\|\textbf{choose}\ ^{\alpha}x\ \textbf{in}\ ^{\beta}v\ \textbf{do}^{\gamma}u\|\!)$	$\rightarrow$	$\mathsf{RemoveEnv}(x)$
		$[\![pos]\!] := (undef, u, undef)$

Choose Rule

$(\!\|\textbf{choose}\ ^{\alpha}x\ \textbf{in}\ ^{\beta}\boxed{e}_1\ \textbf{with}\ ^{\gamma}\boxed{e}_2\ \textbf{do}^{\delta}\boxed{r}\ \|\!) \rightarrow$

$pos := \beta$

$considered(\beta) := \{\}$

$(\!\|\textbf{choose}\ ^{\alpha}x\ \textbf{in}\ ^{\beta}v_1\ \textbf{with}\ ^{\gamma}\boxed{e}_2\ \textbf{do}^{\delta}\boxed{r}\ \|\!) \rightarrow$

if $enumerable(v_1)$ **then**

let $s = enumerate(v_1) \backslash considered(\beta)$ **in**

if $|s| > 0$ **then**

choose $t \in s$ **do**

$\mathsf{AddEnv}(x, t)$

$considered(\beta) := considered(\beta) \cup \{t\}$

$pos := \gamma$

else

$[\![pos]\!] := (undef, \{\!\}, undef)$

else

Error('Cannot choose from non-enumerable element')

$$
\begin{array}{lll}
(\!|\textbf{choose } {}^{\alpha}x \textbf{ in } {}^{\beta}v_1 \textbf{ with } {}^{\gamma}v_2 \textbf{ do}^{\delta}\boxed{\gamma}\ |\!) \rightarrow & \textbf{if } v_2 = \mathsf{true}_e \textbf{ then} \\
 & \quad pos := \delta \\
 & \textbf{else} \\
 & \quad pos := \beta \\
 & \quad \mathsf{RemoveEnv}(x) \\
 & \quad \mathsf{ClearTree}(\gamma) \\
(\!|\textbf{choose } {}^{\alpha}x \textbf{ in } {}^{\beta}v_1 \textbf{ with } {}^{\gamma}v_2 \textbf{ do}^{\delta}u\ |\!) \rightarrow & \mathsf{RemoveEnv}(x) \\
 & [\![pos]\!] := (\mathit{undef}, u, \mathit{undef})
\end{array}
$$

Choose Rule

$$
\begin{array}{lll}
(\!|\textbf{choose } {}^{\alpha}x \textbf{ in } {}^{\beta}\boxed{e}_1 \textbf{ with } {}^{\gamma}\boxed{e}_2 \textbf{ do}^{\delta}\boxed{\gamma}\ \textbf{ ifnone } {}^{\epsilon}\boxed{\gamma}\ |\!) \rightarrow & pos := \beta \\
 & \mathit{considered}(\beta) := \{\} \\
(\!|\textbf{choose } {}^{\alpha}x \textbf{ in } {}^{\beta}v_1 \textbf{ with } {}^{\gamma}\boxed{e}_2 \textbf{ do}^{\delta}\boxed{\gamma}\ \textbf{ ifnone } {}^{\epsilon}\boxed{\gamma}\ |\!) \rightarrow & \\
 & \quad \textbf{if } \mathit{enumerable}(v_1) \textbf{ then} \\
 & \quad\quad \textbf{let } s = \mathit{enumerate}(v_1) \backslash \mathit{considered}(\beta) \textbf{ in} \\
 & \quad\quad\quad \textbf{if } |s| > 0 \textbf{ then} \\
 & \quad\quad\quad\quad \textbf{choose } t \in s \textbf{ do} \\
 & \quad\quad\quad\quad\quad \mathsf{AddEnv}(x, t) \\
 & \quad\quad\quad\quad\quad \mathit{considered}(\beta) := \mathit{considered}(\beta) \cup \{t\} \\
 & \quad\quad\quad\quad pos := \gamma \\
 & \quad\quad\quad \textbf{else} \\
 & \quad\quad\quad\quad pos := \epsilon \\
 & \quad \textbf{else} \\
 & \quad\quad \mathsf{Error}(\text{`Cannot choose from non-enumerable element'}) \\
(\!|\textbf{choose } {}^{\alpha}x \textbf{ in } {}^{\beta}v_1 \textbf{ with } {}^{\gamma}v_2 \textbf{ do}^{\delta}\boxed{\gamma}\ \textbf{ ifnone } {}^{\epsilon}\boxed{\gamma}\ |\!) \rightarrow & \textbf{if } v_2 = \mathsf{true}_e \textbf{ then} \\
 & \quad pos := \delta \\
 & \textbf{else} \\
 & \quad pos := \beta \\
 & \quad \mathsf{RemoveEnv}(x) \\
 & \quad \mathsf{ClearTree}(\gamma) \\
(\!|\textbf{choose } {}^{\alpha}x \in {}^{\beta}v_1 \textbf{ with } {}^{\gamma}v_2 \textbf{ do}^{\delta}u \textbf{ ifnone } {}^{\epsilon}\boxed{\gamma}\ |\!) \rightarrow & \mathsf{RemoveEnv}(x) \\
 & [\![pos]\!] := (\mathit{undef}, u, \mathit{undef}) \\
(\!|\textbf{choose } {}^{\alpha}x \in {}^{\beta}v_1 \textbf{ with } {}^{\gamma}e_2 \textbf{ do}^{\delta}\boxed{\gamma}\ \textbf{ ifnone } {}^{\epsilon}u|\!) \rightarrow & [\![pos]\!] := (\mathit{undef}, u, \mathit{undef})
\end{array}
$$

A.5.2 Forall Rule Plugin

Forall Rule

$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}$□ **do**$^{\gamma}$□ $\|\!)$	$\rightarrow$	$pos := \beta$
		$[\![pos]\!] := (undef, \{\!\}, undef)$
		$considered(\beta) := \{\}$
$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}v$ **do**$^{\gamma}$□ $\|\!)$	$\rightarrow$	**if** $enumerable(v)$ **then**
		let $s = enumerate(v) \backslash considered(\beta)$ **in**
		if $\|s\| > 0$ **then**
		choose $t \in s$ **do**
		AddEnv(x, t)
		$considered(\beta) := considered(\beta) \cup \{t\}$
		$pos := \gamma$
		else
		Error('Cannot enumerate a non-enumerable element')
$(\!\|$ **forall** $^{\alpha}x$ **in** $^{\beta}v$ **do**$^{\gamma}u$ $\|\!)$	$\rightarrow$	$pos := \beta$
		RemoveEnv(x)
		ClearTree(γ)
		$[\![pos]\!] := (undef, updates(pos) \cup u, undef)$

A.5.3 Predicate Logic Plugin

The *and* Operator

Predicate Logic Plugin: and

$(\!\|^{\alpha}$□ **and** $^{\beta}$□$\|\!)_{[400]}$	$\rightarrow$	**choose** $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
		$pos := \lambda$
		ifnone
		if $isBoolean(value(\alpha)) \wedge isBoolean(value(\beta))$ **then**
		if $(value(\alpha) = \mathsf{true}_e) \wedge (value(\beta) = \mathsf{true}_e)$ **then**
		$[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
		else
		$[\![pos]\!] := (undef, undef, \mathsf{false}_e)$

The *or* Operator

Predicate Logic Plugin: or

$$(\!| {}^{\alpha}[?] \textbf{ or } {}^{\beta}[?] |\!)_{[350]} \rightarrow$$

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if $isBoolean(value(\alpha)) \wedge isBoolean(value(\beta))$ **then**
    if $(value(\alpha) = \mathsf{true}_e) \vee (value(\beta) = \mathsf{true}_e)$ **then**
      $[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
    else
      $[\![pos]\!] := (undef, undef, \mathsf{false}_e)$

The *xor* Operator

Predicate Logic Plugin: xor

$$(\!| {}^{\alpha}[?] \textbf{ xor } {}^{\beta}[?] |\!)_{[350]} \rightarrow$$

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if $isBoolean(value(\alpha)) \wedge isBoolean(value(\beta))$ **then**
    if $((value(\alpha) = \mathsf{true}_e) \vee (value(\beta) = \mathsf{true}_e)) \wedge$
        $((value(\alpha) = \mathsf{false}_e) \vee (value(\beta) = \mathsf{false}_e))$ **then**
      $[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
    else
      $[\![pos]\!] := (undef, undef, \mathsf{false}_e)$

The *forall* Universal Quantifier

Predicate Logic Plugin: forall

$(\!\|$ **forall**$^{\alpha}x$ **in** $^{\beta}\boxed{e}$ **holds** $^{\gamma}\boxed{e}$ $\|\!)$	$\rightarrow$	$pos := \beta$
		$considered(\beta) := \{\}$
$(\!\|$ **forall**$^{\alpha}x$ **in** $^{\beta}v$ **holds** $^{\gamma}\boxed{e}$ $\|\!)$	$\rightarrow$	**if** $enumerable(v)$ **then**
		let $s = enumerate(v) \backslash considered(\beta)$ **in**
		if $\|enumerate(v)\| > 0$ **then**
		if $\|s\| > 0$ **then**
		choose $t \in s$ **do**
		AddEnv(x, t)
		$considered(\beta) := considered(\beta) \cup \{t\}$
		$pos := \gamma$
		else
		$[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
		else
		$[\![pos]\!] := (undef, undef, \mathsf{true}_e)$
		else
		Error('Cannot enumerate a non-enumerable element')
$(\!\|$ **forall**$^{\alpha}x$ **in** $^{\beta}v$ **holds** $^{\gamma}v$ $\|\!)$	$\rightarrow$	**if** $(value(\gamma) = \mathsf{true}_e)$ **then**
		$pos := \beta$
		else
		$[\![pos]\!] := (undef, undef, \mathsf{false}_e)$
		RemoveEnv(x)
		ClearTree(γ)

A.5.4 Set Plugin

Set Comprehension Variant 2

Set Plugin : Set Comprehension variant 2

$(\!|\{\ {}^{\alpha}x \mid {}^{\beta_1}x_1 \textbf{ in } {}^{\gamma_1}\boxed{?}_1, \ldots, {}^{\beta_n}x_n \textbf{ in } {}^{\gamma_n}\boxed{?}_n \textbf{ with } {}^{\delta}\boxed{?}\}|\!) \rightarrow$

if $n \geq 1 \wedge \exists j \in [1..n],\ x = x_j$ **then**
  choose $i \in [1..n]$ **with** $\neg evaluated(\gamma_i)$ **do**
    $pos := \gamma_j$
  ifnone
    if $sameNameTwoConstVar$ **then**
      Error('No two constrainer variables may have the same name')
    else if $\exists c \in [1..n], \neg enumerable(value(\gamma_c))$ **then**
      Error('Constrainer variables may only be bound to enumerable elements')
    else if $\exists c \in [1..n], |enumerate(value(\gamma_c))| = 0$ **then**
      $[\![pos]\!] := (undef, undef, newValue(setBack))$
    else
      $newSet(pos) := \{\}$
      InitializeChooseConsideredCombos
      $pos := \delta$
else
  Error('At least one constrainer variable must exist with the same name as the specifier')

where

$sameNameTwoConstVar \equiv \exists k \in [1..n],\ \exists l \in [1..n]\ \ k \neq l \wedge x_k = x_l$

$(\!|\{\ {}^{\alpha}x \mid {}^{\beta_1}x_1 \textbf{ in } {}^{\gamma_1}v_1, \ldots, {}^{\beta_n}x_n \textbf{ in } {}^{\gamma_n}v_n \textbf{ with } {}^{\delta}v\}|\!) \rightarrow$

seq
  if $value(\delta) := \mathsf{true}_e$ **then**
    choose $i \in [1..n]$ **with** $x = x_i$ **do**
      add $env(x_i)$ **to** $newSet(pos)$
next
  if OtherCombosToConsider **then**
    ChooseNextCombo
    ClearTree(δ)
    $pos := \delta$
  else
    DestroyConsideredCombos
    $[\![pos]\!] := (undef, undef, setElement(newSet(pos)))$

Set Comprehension Variant 3

In the following set comprehension form, the guard is optional.

Set Plugin : Set Comprehension variant 3

$(\!|\{\ {}^{\alpha}x \textbf{ is } {}^{\epsilon}\boxed{e}\ \mid {}^{\beta_1}x_1 \textbf{ in } {}^{\gamma_1}\boxed{?}_1, \ldots, {}^{\beta_n}x_n \textbf{ in } {}^{\gamma_n}\boxed{?}_n \textbf{ with } {}^{\delta}\boxed{?}\}|\!) \rightarrow$

if $n \geq 1$ **then**
if $\forall j \in [1..n],\ x \neq x_j$ **then**
choose $j \in [1..n]$ **with** $value(\gamma_j) = undef$ **do**
$pos := \gamma_j$
ifnone
if *sameNameTwoConstVar* **then**
Error('No two constrainer variables may have the same name')
else if $\exists c \in [1..n], \neg enumerable(value(\gamma_c))$ **then**
Error('Constrainer variables may only be bound to enumerable elements')
else if $\exists c \in [1..n], |enumerate(value(\gamma_c))| = 0$ **then**
$[\![pos]\!] := (undef, undef, newValue(setBack))$
else
newSet(pos) $:= \{\}$
InitializeChooseConsideredCombos
$pos := \epsilon$
else
Error('Constrainer variable cannot have same name as specifier')
else
Error('At least one constrainer variable must be present')

where

$sameNameTwoConstVar \equiv \exists k \in [1..n],\ \exists l \in [1..n]\ \ k \neq l \wedge x_k = x_l$

$(\!|\{\ {}^{\alpha}x \textbf{ is } {}^{\epsilon}\boxed{e}\ \mid {}^{\beta_1}x_1 \textbf{ in } {}^{\gamma_1}v_1, \ldots, {}^{\beta_n}x_n \textbf{ in } {}^{\gamma_n}v_n \textbf{ with } {}^{\delta}v\}|\!) \rightarrow$

if $value(\delta) := \mathsf{true}_e$ **then**
$pos := \epsilon$
else
if OtherCombosToConsider **then**
ChooseNextCombo
ClearTree(δ)
$pos := \delta$
else
DestroyConsideredCombos
$[\![pos]\!] := (undef, undef, setElement(newSet(pos)))$

$(\!|\{\ ^{\alpha}x \textbf{ is } ^{\epsilon}v \mid ^{\beta_1}x_1 \textbf{ in } ^{\gamma_1}v_1, \ldots, ^{\beta_n}x_n \textbf{ in } ^{\gamma_n}v_n \textbf{ with } ^{\delta}v\}|\!) \rightarrow$

seq
  add *value*(ϵ) **to** *newSet*(*pos*)
next
  if OtherCombosToConsider **then**
    ChooseNextCombo
    ClearTree(δ)
    ClearTree(ϵ)
    $pos := \delta$
  else
    DestroyConsideredCombos
    $[\![pos]\!] := (undef, undef, setElement(newSet(pos)))$

The Set Difference Operator

Set Plugin : difference

$(\!|^{\alpha}\boxed{?}\backslash^{\beta}\boxed{?}|\!)_{[650]} \rightarrow$

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if SetElement(l) $\wedge$ SetElement(r) **then**
    let $v = \{x \mid x \in enumerate(l) \wedge x \notin enumerate(r)\}$ **in**
      $[\![pos]\!] := (undef, undef, setElement(v))$
where
  $l \equiv value(\alpha), r \equiv value(\beta)$

The Set Union Operator

Set Plugin : union

$(\!|^{\alpha}\boxed{?} \cup ^{\beta}\boxed{?}|\!)_{[650]} \rightarrow$

choose $\lambda \in \{\alpha, \beta\}$ **with** $\neg evaluated(\lambda)$
  $pos := \lambda$
ifnone
  if SetElement(l) $\wedge$ SetElement(r) **then**
    let $v = \{x \mid x \in enumerate(l) \vee x \in enumerate(r)\}$ **in**
      $[\![pos]\!] := (undef, undef, setElement(v))$
where
  $l \equiv value(\alpha), r \equiv value(\beta)$

A.5.5 Math Plugin

Most of the functions provided by the Math plugin are equivalent of their Java counterparts defined in the Java library package `java.lang.Math`. For such functions, we use the descriptions provided by the *Java 2 Platform Standard Edition 5.0 API Specification* [120].

Constants

- `MathE` returns the Number element that is closer in value than any other to e, the base of the natural logarithms.
- `MathPI` returns the Number element that is closer than any other to π, the ratio of the circumference of a circle to its diameter.

Basic Functions

- `abs(v)` returns the absolute value of v.
- `acos(v)` returns the arc cosine of an angle, in the range of 0 through π.
- `asin(v)` returns the arc sine of an angle, in the range of $-\pi/2$ through $\pi/2$.
- `atan(v)` returns the arc tangent of an angle, in the range of $-\pi/2$ through $\pi/2$.
- `atan2(x, y)` converts rectangular coordinates (x, y) to polar (r, θ) and returns θ.
- `cuberoot(v)` returns the cube root of v.
- `cbrt(v)` returns the cube root of v.
- `ceil(v)` returns the smallest (closest to negative infinity) value that is greater than or equal to the argument and is equal to a mathematical integer.
- `cos(v)` returns the trigonometric cosine of an angle.
- `cosh(v)` returns the hyperbolic cosine of v.
- `exp(v)` returns Euler's number e raised to the power of v.
- `expm1(v)` returns $e^v - 1$.

- `floor(v)` returns the largest (closest to positive infinity) value that is less than or equal to the argument and is equal to a mathematical integer.
- `hypot(x, y)` returns $\sqrt{x^2 + y^2}$ without intermediate overflow or underflow.
- `IEEEremainder(v1, v2)` Computes the remainder operation on two arguments as prescribed by the IEEE 754 standard.
- `log(v)` returns the natural logarithm (base e) of v.
- `log10(v)` returns the base 10 logarithm of v.
- `log1p(v)` returns the natural logarithm of the sum of the argument and 1; i.e., $ln(v + 1)$.
- `max(v1, v2)` returns the greater of two values.
- `min(v1, v2)` returns the smaller of two values.
- `pow(x, y)` returns the value of the first argument raised to the power of the second argument.
- `random()` returns a random value with a positive sign, greater than or equal to 0.0 and less than 1.0.
- `round(v)` returns the closest mathematical integer to the argument.
- `signum(v)` Returns zero if the argument is zero, 1.0 if the argument is greater than zero, -1.0 if the argument is less than zero.
- `sin(v)` returns the trigonometric sine of an angle.
- `sinh(v)` returns the hyperbolic sine of v.
- `sqrt(v)` returns the correctly rounded positive square root of v; i.e., $\sqrt{v}$.
- `tan(v)` returns the trigonometric tangent of an angle.
- `tanh(v)` returns the hyperbolic tangent of v.
- `toDegrees(v)` converts an angle measured in radians to an approximately equivalent angle measured in degrees.

- `toRadians(v)` converts an angle measured in degrees to an approximately equivalent angle measured in radians.

Special Functions

- `powerset(set)` computes the powerset of the given set.
- `max({v1,...,vn})` returns the maximum value in a collection of numbers. If there is one non-number in the collection, it returns *undef*.
- `min({v1,...,vn})` returns the minimum value in a collection of numbers. If there is one non-number in the collection, it returns *undef*.
- `sum({v1,...,vn})` returns the sum of a collection of numbers. If there is one non-number in the collection, it returns *undef*.
- `sum({v1,...,vn}, @f)` returns the sum of a collection of numbers, after applying function `f` to the values in the collection. If there is one non-number in the collection, it returns *undef*.
- `powerset({e1,...,en})` returns the powerset of the given set of elements.

Appendix B

CoreASM Examples

B.1 The Railroad Crossing Example

```
CoreASM RailRoadCrossing

use StandardPlugins
use TimePlugin
use MathPlugin

enum Track = {track1, track2}
enum TrackStatus = {empty, coming, crossing}
enum GateSignal = {open, close}
enum GateState = {opened, closed}

function deadline :  Track -> TIME
function trackStatus :  Track -> TrackStatus
function gateSignal :  -> GateSignal
function gateState :  -> GateState

universe Agents = {trackController, gateController, observer, environment}

// Is it safe to open the guard?
  derived safeToOpen = forall t in Track holds
        trackStatus(t) = empty or ( now + dopen) < deadline(t)

derived waitTime = dmin - dclose

init InitRule
```

```
rule InitRule = {
    forall t in Track do {
        trackStatus(t) := empty
        deadline(t) := infinity
    }
    gateState:= opened
    dmin:= 5000
    dmax:= 10000
    dopen:= 2000
    dclose:= 2000
    startTime:= now

    program(trackController) := @TrackControl
    program(gateController) := @GateControl
    program(observer) := @ObserverProgram
    program(environment) := @EnvironmentProgram
    program( self ) := undef
}

rule TrackControl = {
    forall t in Track do {
        SetDeadline(t)
        SignalClose(t)
        ClearDeadline(t)
    }
    SignalOpen
}

rule GateControl = {
    if gateSignal = open and gateState = closed then gateState:= opened
    if gateSignal = close and gateState = opened then gateState:= closed
}

rule SetDeadline(x) =
    if trackStatus(x) = coming and deadline(x) = infinity then
        deadline(x) := now + waitTime

rule SignalClose(x) =
    if now >= deadline(x) and now <= deadline(x) + 1000 then
        gateSignal:= close

rule ClearDeadline(x) =
```

```
        if trackStatus(x) = empty and deadline(x) < infinity then
            deadline(x) := infinity

    rule SignalOpen =
        if gateSignal = close and safeToOpen then
            gateSignal:= open


    // The observer
     rule ObserverProgram =
        seqblock
            print "Time:  " + (( now - startTime) / 1000) + " seconds"
            forall t in Track do
                print "Track " + t + " is " + trackStatus(t)
            print "Gate is " + gateState
            print ""
        endseqblock


    // The environment
     rule EnvironmentProgram =
        choose t in Track do {
            if trackStatus(t) = empty then
                if random < 0.05 then {
                    trackStatus(t) := coming
                    passingTime(t) := now + dmin
                }

            if trackStatus(t) = coming then
                if passingTime(t) < now then {
                    trackStatus(t) := crossing
                    passingTime(t) := now + 4000
                }

            if trackStatus(t) = crossing then
                if passingTime(t) < now then
                    trackStatus(t) := empty
        }
```

B.2 The Surveillance Scenario

```
CoreASM Surveillance_Scenario

use Standard
use Math
use Options

option Signature.NoUndefinedId strict

/* --- Universes --- */
enum Moves = {N, NW, W, WS, S, SE, E, EN}

enum Direction = {forward, away}
universe Agents = {agent1, agent2, environment}

/* --- Function Definitions --- */
// state of the environment


/* --- Function Definitions --- */
// state of the environment
function posX: Agents -> NUMBER
function posY: Agents -> NUMBER
function bearingError:  Agents -> NUMBER
function rangeError:  Agents -> NUMBER

function observationHistory:  Agents -> LIST
function move:Agents -> NUMBER
function dir:  Agents -> Direction

function bearingErrorRange:  Agents -> NUMBER
function rangeErrorRange:  Agents -> NUMBER

// --- Initial Rule ---
init InitRule

rule InitRule = {
    program(agent1) := @Agent1Program
    program(agent2) := @Agent2Program
    program(environment) := @EnvironmentProgram
    program( self ) := undef
```

```
    // initial positions of agents
    posX(agent1) := 0
    posY(agent1) := 0
    posX(agent2) := 15
    posY(agent2) := 10

    dir(agent2) := forward

    // setting error ranges
    bearingErrorRange(agent1) := 3.14 / 20
    rangeErrorRange(agent1) := 2
    bearingErrorRange(agent2) := 3.14 / 20
    rangeErrorRange(agent2) := 4

    // initial values of agent functions
    forall a in {agent1, agent2} do {
        observationHistory(a) := []
        bearingError(a) := 0
        rangeError(a) := 0
    }
}

// --- Agent Programs ---
 rule Agent1Program = {
    RecordObservation(agent2)
    if isInAOI(agent2) then
        SendMessage( "Agent 2 is in the area of interest." )
    if size(observationHistory( self )) > 1 then
        if approaching( self ) then
            print "Agent 1:  Agent 2 is approaching."
}

rule Agent2Program =  {
    RecordObservation(agent1)
    if dir( self ) = forward then
        MoveToward(agent1)
    else
        MoveAwayFrom(agent1)

    if tooClose(agent1) then
        dir( self ) := away
}
```

```
rule EnvironmentProgram =
    forall a in {agent1, agent2} do {
        bearingError(a) := bearingErrorRange(a) * (2 * random - 1)
        rangeError(a) := rangeErrorRange(a) * (2 * random - 1)
    }

// --- Auxiliary Rules ---
 rule RecordObservation(a) =
    add [obsRange(self, a), obsBearing(self, a)] to observationHistory(self)

rule SendMessage(msg) =
    "SendMessage(" + msg + ")"

rule Move(dir) = {
    print "agent1:(" + posX(agent1) + ", " + posY(agent1)
            + ") - agent2:(" + posX(agent2) + ", " + posY(agent2) + ")"
    if dir = N then
        posY( self ) := posY( self ) + 1
    else if dir = S then
        posY( self ) := posY( self ) - 1
    else if dir = W then
        posX( self ) := posX( self ) - 1
    else if dir = E then
        posX( self ) := posX( self ) + 1
    else if dir = EN then {
        Move(N)
        Move(E)
    }
    else if dir = NW then {
        Move(N)
        Move(W)
    }
    else if dir = SE then {
        Move(S)
        Move(E)
    }
    else if dir = WS then {
        Move(S)
        Move(W)
    }
}
```

```
/* Move towards agent 'a' */
 rule MoveToward(a) =
    let dir = getDirection(
                atan2(posX(agent1) - posX(self), posY(agent1) - posY(self))
                + (2 * random * bearingError(self) - bearingError(self))
              ) in
        Move(dir)

/* Move away from agent 'a' */
 rule MoveAwayFrom(a) =
    let nb = atan2(posX(agent1) - posX(self), posY(agent1) - posY(self))
                + (2 * random * bearingError(self) - bearingError(self))
                - signum(atan2(posX(agent1) - posX(self),
                        posY(agent1) - posY(self))
                    + (2 * random * bearingError(self) - bearingError(self)))
                * MathPI in
        Move(getDirection(nb))

// Compute a move direction based on the given bearing
 rule getDirection(b) =
    return move in
        let bp = abs(b) in {
            if bp < ( MathPI / 8) then
                move:= N
            if abs(bp - MathPI / 4) < ( MathPI / 8) then
                if (b < 0) then
                    move:= EN
                else
                    move:= NW
            if abs(bp - MathPI / 2) < ( MathPI / 8) then
                if (b < 0) then
                    move:= E
                else
                    move:= W
            if abs(bp - (3 * MathPI / 4)) < ( MathPI / 8) then
                if (b < 0) then
                    move:= WS
                else
                    move:= SE
            if abs(bp - MathPI) < ( MathPI / 8) then
                move:= S
```

```
        }

    /* ----- Derived Functions ----- */

    derived bearing(a) = atan2(posX(a) - posX( self ), posY(a) - posY( self ))

    derived range(a) =
        sqrt( pow(posX(a) - posX( self ), 2) + pow(posY(a) - posY( self ), 2))

    derived obsBearing(observer, observed) =
        bearing(observed) + bearingError(observer)

    derived obsRange(observer, observed) =
        range(observed) + rangeError(observer)

    derived isInAOI(a) =
        obsRange( self , a) > 5 and obsRange( self , a) < 12
        and obsBearing( self , a) < ( MathPI / 3)
        and obsBearing( self , a) > ( MathPI / 6)

    derived tooClose(observed) =
        obsRange( self , observed) < 12

    derived approaching(observer) =
        head( last(observationHistory(observer)))
        < head( nth(observationHistory(observer),
                    size(observationHistory(observer)) - 1))
```

Bibliography

[1] *The Object Management Group (OMG)*. http://www.omg.org.

[2] J.R. Abrial. *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, 1996.

[3] M. Altenhofen, A. Friesen, and J. Lemcke. Asms in service oriented architectures. *Journal of Universal Computer Science*, 14(12):2034–2058, 2008.

[4] M. Anlauff. XASM – An Extensible, Component-Based Abstract State Machines Language. In Y. Gurevich and P. Kutter and M. Odersky and L. Thiele, editor, *Abstract State Machines: Theory and Applications*, volume 1912 of *LNCS*, pages 69–90. Springer-Verlag, 2000.

[5] M. Anlauff and P. Kutter. *eXtensible Abstract State Machines*. XASM open source project: http://www.xasm.org.

[6] Jörg Beckers, Daniel Klünder, Stefan Kowalewski, and Bastian Schlich. Direct support for model checking abstract state machines by utilizing simulation. In *ABZ '08: Proceedings of the 1st international conference on Abstract State Machines, B and Z*, pages 112–124, Berlin, Heidelberg, 2008. Springer-Verlag.

[7] B. Beckert and J. Posegga. leanEA: A Lean Evolving Algebra Compiler. In H. Kleine Büning, editor, Proceedings of the Annual Conference of the European Association for Computer Science Logic (CSL'95), volume 1092 of *LNCS*, pages 64–85. Springer, 1996.

[8] C. Beierle, E. Börger, I. Durdanovic, U. Glässer, and E. Riccobene. Refining Abstract Machine Specifications of the Steam Boiler Control to Well Documented Executable Code. In J.-R. Abrial, E. Börger, and H. Langmaack, editors, *Formal Methods for Industrial Applications. Specifying and Programming the Steam-Boiler Control*, number 1165 in LNCS, pages 62–78. Springer, 1996.

[9] Daniel M. Berry. Formal Methods: the very idea—Some thoughts about why they work when they work. *Science of Computer Programming*, 42(1):11–27, 2002.

[10] M. Bidoit and Peter Mosses. *Casl User Manual: Introduction to Using the Common Algebraic Specification Language Casl*. SpringerVerlag, 2004.

[11] Dines Bjørner and Cliff B. Jones, editors. *The Vienna Development Method: The Meta-Language*, volume 61 of *Lecture Notes in Computer Science*. Springer, 1978.

[12] A. Blass and Y. Gurevich. Background, Reserve, and Gandy Machines. In P. Clote and H. Schwichtenberg, editors, *Computer Science Logic (Proceedings of CSL 2000)*, volume 1862 of *LNCS*, pages 1–17. Springer, 2000.

[13] Andreas Blass and Yuri Gurevich. Abstract State Machines Capture Parallel Algorithms. *ACM Transactions on Computation Logic*, 4(4):578–651, 2003.

[14] E. Börger. A Logical Operational Semantics for Full Prolog. Part I: Selection Core and Control. In E. Börger, H. Kleine Büning, M. M. Richter, and W. Schönfeld, editors, *CSL'89. 3rd Workshop on Computer Science Logic*, volume 440 of *LNCS*, pages 36–64. Springer, 1990.

[15] E. Börger. A Logical Operational Semantics of Full Prolog. Part II: Built-in Predicates for Database Manipulation. In B. Rovan, editor, *Mathematical Foundations of Computer Science*, volume 452 of *LNCS*, pages 1–14. Springer, 1990.

[16] E. Börger. Computation and Specification Models: A Comparative Study. In P. Mosses, editor, *Proceedings of the Fourth International Workshop on Action Semantics, AS 2002*, number NS-02-8 in BRICS Notes Series, pages 110–133. University of Aarhus, Department of Computer Science, 2002.

[17] E. Börger. The ASM ground model method as a foundation of requirements engineering. In N.Dershowitz, editor, *Verification: Theory and Practice*, volume 2772 of *LNCS*, pages 145–160. Springer-Verlag, 2003.

[18] E. Börger. The ASM refinement method. *Formal Aspects of Computing*, 15:237–257, 2003.

[19] E. Börger, N. G. Fruja, V. Gervasi, and R. F. Stärk. A High-level Modular Definition of the Semantics of C#. *Theoretical Computer Science*, 336(2/3):235–284, May 2005.

[20] E. Börger, U. Glässer, and W. Müller. The Semantics of Behavioral VHDL'93 Descriptions. In *EURO-DAC'94. European Design Automation Conference with EURO-VHDL'94*, pages 500–505, Los Alamitos, California, 1994. IEEE CS Press.

[21] E. Börger, U. Glässer, and W. Müller. Formal Definition of an Abstract VHDL'93 Simulator by EA-Machines. In C. Delgado Kloos and P. T. Breuer, editors, *Formal Semantics for VHDL*, pages 107–139. Kluwer Academic Publishers, 1995.

[22] E. Börger, P. Päppinghaus, and J. Schmid. Report on a Practical Application of ASMs in Software Design. In Y. Gurevich and P. Kutter and M. Odersky and L. Thiele, editor, *Abstract State Machines: Theory and Applications*, volume 1912 of *LNCS*, pages 361–366. Springer-Verlag, 2000.

[23] E. Börger, E. Riccobene, and J. Schmid. Capturing Requirements by Abstract State Machines: The Light Control Case Study. *Journal of Universal Computer Science*, 6(7):597–620, 2000.

[24] E. Börger and W. Schulte. A Practical Method for Specification and Analysis of Exception Handling: A Java/JVM Case Study. *IEEE Transactions on Software Engineering*, 26(10):872–887, October 2000.

[25] E. Börger and R. Stärk. *Abstract State Machines: A Method for High-Level System Design and Analysis*. Springer-Verlag, 2003.

[26] Egon Börger. Construction and Analysis of Ground Models and their Refinements as a Foundation for Validating Computer Based Systems. *Formal Aspects of Computing*, 19(2):225–241, 2007.

[27] É. Bossé, A.-L. Jousselme, and P. Maupin. Situation Analysis for Decision Support: A Formal Approach. In *Proc. of the 10th Intl. Conf. on Information Fusion*, July 2007.

[28] É. Bossé, J. Roy, and S. Ward. *Models and Tools for Information Fusion*. 2007.

[29] P. J. Brantingham and P. L. Brantingham. *Patterns in Crime*. New York: Macmillan Publishing Company, 1984.

[30] P. L. Brantingham, U. Glässer, P. Jackson, and M. Vajihollahi. Modeling Criminal Activity in Urban Landscapes. Technical Report SFU-CMPT-TR-2008-13, Simon Fraser University, Aug 2008.

[31] P. L. Brantingham, U. Glässer, B. Kinney, K. Singh, and M. Vajihollahi. A Computational Model for Simulating Spatial Aspects of Crime in Urban Environments. In M. Jamshidi, editor, *Proceedings of the 2005 IEEE International Conference on Systems, Man and Cybernetics*, pages 3667–74, October 2005.

[32] P. L. Brantingham, B. Kinney, U. Glässer, P. Jackson, and M. Vajihollahi. Mastermind: Computational Modeling and Simulation of Spatiotemporal Aspects of Crime in Urban Environments. In L. Liu and J. Eck, editors, *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*. Information Science Reference, 2008.

[33] Colin Campbell, Wolfgang Grieskamp, Lev Nachmanson, Wolfram Schulte, Nikolai Tillmann, and Margus Veanes. Model-Based Testing of Object-Oriented Reactive Systems with Spec Explorer. Technical Report MSR-TR-2005-59, Microsoft FSE Group, May 2005.

[34] Ora Canada. Z/eves version 1.5: An overview. In *FM-Trends*, pages 367–376, 1998.

[35] Giuseppe Del Castillo. The ASM Workbench: an Open and Extensible Tool Environment for Abstract State Machines. In *Workshop on Abstract State Machines*, pages 139–154, 1998.

[36] Giuseppe Del Castillo. *The ASM Workbench: A Tool Environment for Computer-Aided Analysis and Validation of Abstract State Machine Models*. PhD thesis, Informatik und Heinz Nixdorf Institut, Universität Paderborn, Germany, 2000.

[37] A. Cavarra and E. Riccobene. Simulating UML Statecharts. In R. Moreno-Díaz and A. Quesada-Arencibia, editors, *Formal Methods and Tools for Computer Science (Proceedings of Eurocast 2001)*, pages 224–227, Canary Islands, Spain, February 2001. Universidad de Las Palmas de Gran Canaria.

[38] G. Del Castillo. Towards Comprehensive Tool Support for Abstract State Machines. In D. Hutter, W. Stephan, P. Traverso, and M. Ullmann, editors, *Applied Formal Methods — FM-Trends 98*, volume 1641 of *LNCS*, pages 311–325. Springer-Verlag, 1999.

[39] G. Del Castillo, I. Durdanović, and U. Glässer. An Evolving Algebra Abstract Machine. In H. Kleine Büning, editor, Proceedings of the Annual Conference of the European Association for Computer Science Logic (CSL'95), volume 1092 of *LNCS*, pages 191–214. Springer, 1996.

[40] Matteo Demuru. Modeling cell methabolic mechanisms through Abstract State Machines. Master's thesis, University of Pisa, Italy, February 2008.

[41] D. Diesen. *Specifying Algorithms Using Evolving Algebra. Implementation of Functional Programming Languages*. Dr. scient. degree thesis, Dept. of Informatics, University of Oslo, Norway, March 1995.

[42] Jan Ellsberger, Dieter Hogrefe, and Amardeo Sarma. *SDL : Formal Object-oriented Language for Communicating Systems*. Prentice Hall, 1997.

[43] M. R. Endsley. Theoretical Underpinnings of Situation Awareness: A Critical Review. In M. R. Endsley and D. J. Garland, editors, *Situation Awareness Analysis and Measurement*. LEA, 2000.

[44] R. Eschbach, U. Gässer, R. Gotzhein, and A. Prinz. On the Formal Semantics of SDL-2000: A Compilation Approach Based on an Abstract SDL Machine. In Y. Gurevich and P. Kutter and M. Odersky and L. Thiele, editor, *Abstract State Machines: Theory and Applications*, volume 1912 of *LNCS*, pages 242–265. Springer-Verlag, 2000.

[45] R. Eschbach, U. Glässer, R. Gotzhein, M. von Löwis, and A. Prinz. Formal Definition of SDL-2000: Compiling and Running SDL Specifications as ASM Models. *Journal of Universal Computer Science*, 7(11):1024–1049, 2001.

[46] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, 2003.

[47] R. Farahbod, V. Gervasi, and U. Glässer. Design and Specification of the CoreASM Execution Engine. Technical Report SFU-CMPT-TR-2005-02, Simon Fraser University, February 2005.

[48] R. Farahbod, V. Gervasi, and U. Glässer. CoreASM: An Extensible ASM Execution Engine. *Fundamenta Informaticae*, pages 71–103, 2007.

[49] R. Farahbod and U. Glässer. Semantic Blueprints of Discrete Dynamic Systems: Challenges and Needs in Computational Modeling of Complex Behavior. In *New Trends in Parallel and Distributed Computing, Proc. 6th Intl. Heinz Nixdorf Symposium, Jan. 2006*, pages 81–95. Heinz Nixdorf Institute, 2006.

[50] R. Farahbod, U. Glässer, É. Bossé, and A. Guitouni. Integrating Abstract State Machines and Interpreted Systems for Situation Analysis Decision Support Design. In *Proc. of the 11th Intl Conf. on Information Fusion (Fusion 2008)*, July 2008.

[51] R. Farahbod, U. Glässer, P. Jackson, and M. Vajihollahi. High Level Analysis, Design and Validation of Distributed Mobile Systems with CoreASM. In *Proceedings of 3rd International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (ISoLA 2008)*. Springer, October 2008.

[52] R. Farahbod, U. Glässer, and A. Khalili. A Multi-Layer Network Architecture for Dynamic Resource Configuration & Management of Multiple Mobile Resources in Maritime Surveillance. In *Proc. of SPIE Defense & Security Symposium*, March 2009.

[53] R. Farahbod, U. Glässer, and M. Vajihollahi. Specification and Validation of the Business Process Execution Language for Web Services. In Wolf Zimmermann and Bernhard Thalheim, editors, *Abstract State Machines 2004. Advances In Theory And Practice: 11th International Workshop (ASM 2004)*, Germany, March 2004. Springer-Verlag.

[54] R. Farahbod, U. Glässer, and M. Vajihollahi. A Formal Semantics for the Business Process Execution Language for Web Services. In Savitri Bevinakoppa et al., editors, *Web Services and Model-Driven Enterprise Information Systems*, pages 144—155, Portugal, May 2005. INSTICC Press.

[55] R. Farahbod, U. Glässer, and M. Vajihollahi. Abstract Operational Semantics of the Business Process Execution Language for Web Services. Technical Report SFU-CMPT-TR-2005-04, Simon Fraser University, Feb. 2005. Revised version of SFU-CMPT-TR-2004-03, April 2004.

[56] R. Farahbod, U. Glässer, and M. Vajihollahi. An Abstract Machine Architecture for Web Service Based Business Process Management. *International Journal of Business Process Integration and Management*, 1:279–291, 2007.

[57] R. Farahbod, U. Glässer, and H. Wehn. CanCoastWatch Dynamic Configuration Manager. In *Proc. of the 14th Intl. Abstract State Machines Workshop*, June 2007.

[58] R. Farahbod, U. Glässer, and H. Wehn. Dynamic Resource Management for Adaptive Distributed Information Fusion in Large Volume Surveillance. In *Proc. of SPIE Defense & Security Symposium*, March 2008.

[59] R. Farahbod, Uwe Glässer, and G. Ma. Model Checking CoreASM Specifications. In A. Prinz, editor, *Proceedings of the 14th International ASM Workshop (ASM'07)*, 2007.

[60] Formal Methods laboratory of University of Milan. *Asmeta*, 2006. Last visited June 2008, http://asmeta.sourceforge.net/.

[61] Free Software Foundation. *GNU General Public License*, 2007. Available electronically at http://www.gnu.org/copyleft/gpl.html (Last visited in March 2009).

[62] Free Software Foundation. *GNU Lesser General Public License*, 2007. Available electronically at http://www.gnu.org/copyleft/lgpl.html (Last visited in March 2009).

[63] The Apache Software Foundation. *Apache License*, 2004. Available electronically at http://www.apache.org/licenses (Last visited in March 2009).

[64] Martin Fowler. The New Methodology. April 2003. http://martinfowler.com/articles/newMethodology.html.

[65] R. France, A. Evans, K. Lano, and B. Rumpe. The UML as a formal modeling notation. *Comput. Stand. Interfaces*, 19(7):325–334, 1998.

[66] Angelo Gargantini, Elvinia Riccobene, and Salvatore Rinzivillo. Using Spin to Generate Tests from ASM Specifications. In *Abstract State Machines 2003*, pages 263–277. Springer, 2003.

[67] Angelo Gargantini, Elvinia Riccobene, and Patrizia Scandurra. A Metamodel-based Simulator for ASMs. In *Proc. of the 14th Intl. Abstract State Machines Workshop*, June 2007.

[68] V. Gervasi and R. Farahbod. JASMine: Accessing java code from CoreASM. In *Proceedings of the Dagstuhl Seminar on Rigorous Methods for Software Construction and Analysis (LNCS Festschrift)*. Springer, 2009 (to be published).

[69] U. Glässer, R. Gotzhein, and A. Prinz. The Formal Semantics of SDL-2000: Status and Perspectives. *Computer Networks*, 42(3):343–358, 2003.

[70] U. Glässer and Q.-P. Gu. Formal Description and Analysis of a Distributed Location Service for Mobile Ad Hoc Networks. *Theoretical Comp. Sci.*, 336:285–309, May 2005.

[71] U. Glässer, Y. Gurevich, and M. Veanes. Abstract Communication Model for Distributed Systems. *IEEE Trans. on Soft. Eng.*, 30(7):458–472, July 2004.

[72] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java Language Specification*. Prentice Hall, third edition, 2005.

[73] Y. Gurevich. Evolving Algebras. A Tutorial Introduction. *Bulletin of EATCS*, 43:264–284, 1991.

[74] Y. Gurevich. Evolving Algebras 1993: Lipari Guide. In E. Börger, editor, *Specification and Validation Methods*, pages 9–36. Oxford University Press, 1995.

[75] Y. Gurevich and J. Huggins. Evolving Algebras and Partial Evaluation. In B. Pehrson and I. Simon, editors, *IFIP 13th World Computer Congress*, volume I: Technology/Foundations, pages 587–592, Elsevier, Amsterdam, the Netherlands, 1994.

[76] Y. Gurevich and N. Tillmann. Partial Updates: Exploration. *Journal of Universal Computer Science*, 7(11):917–951, 2001.

[77] Y. Gurevich and N. Tillmann. Partial Updates. *Journal of Theoretical Computer Science*, 336(2-3):311–342, 2005.

[78] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969.

[79] C. A. R. Hoare. Communicating sequential processes. pages 413–443, 2002.

[80] Gerard J. Holzmann. The Model Checker SPIN. *IEEE Trans. Software Eng.*, 23(5):279–295, 1997.

[81] J. Huggins. An offline partial evaluator for evolving algebras. Technical Report CSE-TR-229-95, University of Michigan, 1995.

[82] J. Huggins and C. Wallace. An Abstract State Machine Primer. Technical Report CS-TR-02-04, Computer Science Department, Michigan Technological University, 4 December 2002.

[83] ITU-T Recommendation Z.100 Annex F (11/00). *SDL Formal Semantics Definition*. International Telecommunication Union, 2001.

[84] Daniel Jackson. *Software Abstractions: Logic, Language, and Analysis*. MIT Press, 2006.

[85] Olav Jensen, Raymond Koteng, Kjetil Monge, and Andreas Prinz. Abstraction using ASM Tools. In A. Prinz, editor, *Proceedings of the 14th International ASM Workshop (ASM'07)*, 2007.

[86] C. W. Johnson. Literate specifications. *Software Engineering Journal*, 11(4):225–237, July 1996.

[87] A. M. Kappel. Executable Specifications Based on Dynamic Algebras. In A. Voronkov, editor, *Logic Programming and Automated Reasoning*, volume 698 of *Lecture Notes in Artificial Intelligence*, pages 229–240. Springer, 1993.

[88] K.L. McMillan. The SMV system. Technical Report CMU-CS-92-131, 1992.

[89] Donald E. Knuth. Literate programming. *Comput. J.*, 27(2):97–111, 1984.

[90] William Leiserson. *Elegant, efficient LL (k) parser generation.* PhD thesis, Rochester Institute of Technology, Rochester, USA, 2006.

[91] Jens Lemcke and Andreas Friesen. Composing web-service-like abstract state machines (asms). *Services, IEEE Congress on*, pages 262–269, 2007.

[92] P. Lucas and K. Walk. On the formal description of PL/I. *Annual Review of Automatic Programming*, 6:105–182, 1969.

[93] George Z. Ma. Model Checking Support for CoreASM: Model Checking Distributed Abstract State Machines Using Spin. Master's thesis, Simon Fraser University, Canada, May 2007.

[94] Patrick Maupin and Anne-Laure Jousselme. A General Algebraic Framework for Situation Analysis. In *Proc. of the 8^{th} Intl. Conf. on Information Fusion*, Philadelphia, PA, July 2005.

[95] Patrick Maupin and Anne-Laure Jousselme. Interpreted Systems for Situation Analysis. In *Proc. of the 10th Intl. Conf. on Information Fusion*, Quebec city, Canada, 9-12 July 2007.

[96] M. Mauve, J. Widmer, and H. Hartenstein. A survey on position-based routing in mobile ad-hoc networks. *IEEE Network*, 15, 2001.

[97] Daniele Mazzei, Federico Vozzi, Antonio Cisternino, Giovanni Vozzi, and Arti Ahluwalia. A high-throughput bioreactor system for simulating physiological environment. *IEEE Transactions on Industrial Electronics*, 55(9):3273–3280, 2008.

[98] Stephen J. Mellor, Kendall Scott, Axel Uhl, and Dirk Weise. *MDA Distilled: Principles of Model-Driven Architecture.* Addison-Wesley, 2004.

[99] Mashaal A. Memon. Specification language design concepts: Aggregation and extensibility in coreasm. Master's thesis, Simon Fraser University, Burnaby, Canada, April 2006.

[100] Microsoft Corp. *Microsoft .NET Framework.* Last visited Dec. 2006, `http://www.microsoft.com/net`.

[101] Microsoft FSE Group. *The Abstract State Machine Language*, 2003. Last visited June 2008, `http://research.microsoft.com/fse/asml/`.

[102] Microsoft FSE Group. *AsmL Community Project on CodePlex*, 2008. Last visited July 2008, `http://www.codeplex.com/AsmL`.

[103] Microsoft FSE Group. *Spec Explorer*, 2008. Last visited July 2008, `http://research.microsoft.com/specexplorer`.

[104] R. Milner, M. Tofte, and R. Harper. *The Definition of Standard ML.* MIT Press, 1990.

[105] Robin Milner, Joachim Parrow, and David Walker. A Calculus of Mobile Processes. *Information and Computation*, 100:1–40, September 1992.

[106] Peter D. Mosses. *Action semantics*. Cambridge University Press, New York, NY, USA, 1992.

[107] Peter D. Mosses. A Tutorial on Action Semantics. Technical Report NS-96-14, Basic Research in Computer Science (BRICS), 1996.

[108] W. Müller, J. Ruf, and W. Rosenstiel. An ASM Based SystemC Simulation Semantics. In W. Müller et al., editors, *SystemC - Methodologies and Applications*. Kluwer Academic Publishers, June 2003.

[109] Regents of the University of California. *BSD Licenses*, 1990-2009. Available electronically at `http://en.wikipedia.org/wiki/BSD_licenses` (Last visited in March 2009).

[110] Frank G. Pagan. *Formal Specification of Programming Languages: A Panoramic Primer*. Prentice Hall, February 1981.

[111] J. L. Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice-Hall, 1981.

[112] Joachim Schmid. *Introduction to AsmGofeer*, March 2001. Available electronically at `http://www.tydo.de/Doktorarbeit/AsmGofer/files/AsmGoferIntro.pdf` (Last visited in July 2008).

[113] Joachim Schmid. *AsmGofer*, 2008. Available electronically at `http://www.tydo.de/Doktorarbeit/AsmGofer/` (Last visited in July 2008).

[114] Douglas C. Schmidt. Model-Driven Engineering. *IEEE Computer*, 39, February 2006.

[115] Dana Scott and Christopher Strachey. [toward a mathematical semantics for computer languages. Technical report.

[116] J. Michael Spivey. *The Z Notation: a reference manual*. Prentice Hall International Series in Computer Science, 2 edition, 1992.

[117] Thomas A. Standish. Extensibility in programming language design. *SIGPLAN Not.*, 10(7):18–21, 1975.

[118] R. Stärk, J. Schmid, and E. Börger. *Java and the Java Virtual Machine: Definition, Verification, Validation*. Springer-Verlag, 2001.

[119] Joseph E. Stoy. *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*. The MIT Press, September 1981.

[120] Sun Microsystems, Inc. *The Java 2 Platform Standard Edition 5.0 API Specification*. Sun Microsystems, Inc., 2004. (`http://java.sun.com/j2se/1.5.0/docs/api`).

[121] Clemens Szyperski. *Component Software: Beyond Object-Oriented Programming*. Addison-Wesley Professional, December 1997.

[122] Margus Veanes, Colin Campbell, Wolfgang Grieskamp, Wolfram Schulte, Nikolai Tillmann, and Lev Nachmanson. Model-Based Testing of Object-Oriented Reactive Systems with Spec Explorer. In Robert M. Hierons, Jonathan P. Bowen, and Mark Harman, editors, *Formal Methods and Testing*, volume 4949 of *Lecture Notes in Computer Science*, pages 39–76. Springer, 2008.

[123] H. Wehn et al. A Distributed Information Fusion Testbed for Coastal Surveillance. In *Proc. of the 10th Intl. Conf. on Information Fusion*, July 2007.