

A Reachability Analyser for Algebraic System Nets

Marko Mäkelä

March 2000

Licentiate's Thesis

Helsinki University of Technology
Department of Computer Science and Engineering
Laboratory for Theoretical Computer Science

Tekijä: Marko Mäkelä Työn nimi: A Reachability Analyser for Algebraic System Nets suomeksi: Saavutettavuusanalysointia algebrallisille järjestelmäverkoille	
Päivämäärä: 27. maaliskuuta 2000	Sivumäärä: 96
Osasto: Tietotekniikka Professuuri: Tik-79 Digitaalitekniikka	
Valvoja: Prof. Leo Ojala	Ohjaaja: Prof. Nisse Husberg
Tiivistelmä: <p>Rinnakkaisia ja hajautettuja järjestelmiä on vaikea hallita käyttämättä formaaleja ja analysointimenetelmiä. Menetelmien soveltajan on tasapainoitettava ilmaisuvoimaisuuden ja hallittavuuden välillä. Ilmaisuvoimaltaan heikot menetelmät eivät sovellu kaikkien todellisten järjestelmien kuvaamiseen, mutta kattavan saavutettavuusanalyysin tekeminen erittäin ilmaisuvoimaisella menetelmällä määritellyn järjestelmän tiloista voi olla ratkeamaton ongelma.</p> <p>Tämä työ määrittelee tekijän toteuttaman, tietokoneohjelmistopohjaisten yhteyskäytäntöjen analysointiin suunnitellun saavutettavuusanalysointia algebrallisille monilajialgebrojen avulla. Työ esittelee myös analysointia algebrallisilla mallinnusformalismin, algebralliset järjestelmäverkot.</p> <p>Työ kuvaa saavutettavuusanalysointia algebrallisilla toteutusyksityiskohtia sekä teoreettisesta että ohjelmistoteknisestä näkökulmasta. Lopuksi työssä näytetään, kuinka eräitä muille työkaluille hankalia rakenteita voidaan kuvata esitetyllä menetelmällä ja kuinka erilaisia ohjelmointikielten rakenteita voidaan puoli-automaattisesti kääntää analysoitavan mallin osiksi.</p> Avainsanat: monilajialgebrat, algebralliset järjestelmäverkot, hajautetut järjestelmät, saavutettavuusanalyysi	

Author: Marko Mäkelä Name of the Thesis: A Reachability Analyser for Algebraic System Nets in Finnish: Saavutettavuusanalyysaattori algebrallisille järjestelmäverkoille	
Date: 27th March 2000	Pages: 96
Department: Computer Science and Engineering Professorship: Tik-79 Digital Systems Science	
Supervisor: Prof. Leo Ojala	Instructor: Prof. Nisse Husberg
Abstract: <p>Concurrent and distributed systems are difficult to manage without using formal analysis methods. The user of formal methods has to find a balance between expressive power and tractability. A formalism with small expressive power may not suit well to describing all real-life systems, but on the other hand, performing exhaustive reachability analysis on a system defined using a highly expressive formalism may be an unsolvable problem.</p> <p>Using extended many-sorted algebras, this work defines the data type system and the set of algebraic operations that the author has implemented in a reachability analyser intended for modelling computer software based communications protocols. The thesis also introduces the modelling formalism of the analyser, algebraic system nets.</p> <p>The thesis discusses some implementation details of the reachability analyser both from a theoretical and from a software technology point of view. Finally, the work shows how some constructs difficult for other tools can be modelled using the new formalism, and describes how different programming language constructs can be transformed to parts of a model by a semi-automated compiler.</p> Keywords: many-sorted algebras, algebraic system nets, distributed systems, reachability analysis	

Preface

This research forms my Licentiate of Technology thesis. I wish to thank Professor Nisse Husberg, Kimmo Varpaaniemi, D.Sc. (Tech.), Tommi Junttila, Keijo Heljanko and Tommi Syrjänen from the Laboratory for Theoretical Computer Science of Helsinki University of Technology for discussions, feedback and numerous valuable ideas. Also, I wish to thank Tommi Junttila for introducing me to Algebraic System Nets, the formalism the formal notations of this work are based on. Special thanks go to Professor Leo Ojala, the supervisor of the thesis, who encouraged me to carry out the work further than I originally had planned.

The research took place in the MARIA project financed by the National Technology Agency of Finland (TEKES), Nokia Research Center, Nokia Networks, the Helsinki Telephone Corporation and the Finnish Rail Administration. The work was also supported by the Helsinki Graduate School in Computer Science and Engineering (HeCSE) and by a personal grant from Tekniikan Edistämissäätiö.

I would like to thank my friends for all kinds of support I have received during this research project, and Mikko Suonio in particular for reviewing this work and for his constructive comments.

Marko Mäkelä
Otaniemi, 27th March 2000

Contents

I	Identifying the Problems	1
1	Introduction	2
1.1	Reachability Analysers in the Past and Present	4
1.2	Related Work	6
1.3	Outline	6
2	Verification of Concurrent Programs	7
2.1	Classifying Programs	7
2.1.1	Sequential Behaviour	7
2.1.2	Concurrent Behaviour	8
2.2	Making Abstractions	8
2.2.1	Atomising Sequences of Actions	8
2.2.2	Introducing Nondeterminism	9
3	Computer Tools for Analysis	12
3.1	Incomplete Methods	12
3.1.1	Static Analysis	12
3.1.2	Instrumenting Program Code	13
3.1.3	Regression Testing	14
3.2	Formal Methods	14
3.2.1	Constructing the Model	15
3.2.2	Analysing the Model	17
II	Designing the Formalism	19
4	Algebraic System Nets	20
4.1	Signatures and Algebras	20
4.1.1	Signatures, Variables and Terms	21
4.1.2	Algebras, Assignments and Evaluations	22
4.1.3	Multi-Set Signatures and Algebras	23
4.2	Algebraic System Nets	24
5	Data Types	29
5.1	Design Criteria	29
5.1.1	Tight Representation	30

5.1.2	Expressive Power	30
5.2	Simple Types	30
5.2.1	Boolean	31
5.2.2	Character	31
5.2.3	Integer	31
5.2.4	Enumerated Types	31
5.2.5	Identifier Type	32
5.3	Structured Types	32
5.3.1	Tuple	32
5.3.2	Associative Array	33
5.3.3	Variable-Length Buffer	33
5.3.4	Tagged Union	34
5.4	Constraints	34
5.4.1	Computing the Union	35
5.4.2	Computing the Intersection	35
6	Algebraic Operations	36
6.1	Design Criteria	37
6.2	Variables	37
6.3	Operations on Basic Sorts	37
6.3.1	Constants	37
6.3.2	Total Order	38
6.3.3	Logical Operations	39
6.3.4	Integer Arithmetics	40
6.3.5	Structure Operations	41
6.3.6	Type Conversions	43
6.4	Operations on Multi-Set Sorts	44
6.4.1	Multi-Set Constructor	44
6.4.2	Empty Multi-Set	45
6.4.3	Multi-Set Sum and Filter	45
6.4.4	Multi-Set Transformations	46
6.4.5	Union and Intersection	46
6.4.6	Scalar Multiplication	47
6.4.7	Comparison	47
6.4.8	Cardinality	47
6.5	Short-Circuit Operations	47
III	Implementing and Using the Analyser	49
7	Implementing the Analyser	50
7.1	Transition Instance Analysis	50
7.1.1	Rough Description of the Algorithm	53
7.1.2	Unfolding to Place/Transition Nets	55
7.2	The Expression Evaluator	56
7.2.1	Error Handling	56

7.2.2	Optimisations	57
7.2.3	Interpreting vs. Compiling	58
7.3	Managing the Reachability Graph	61
7.3.1	Encoding Markings	63
7.3.2	Encoding Transition Instances	65
7.4	MARIA as a Software Project	66
7.4.1	Development Tools	66
7.4.2	Selected Implementation Details	69
8	Constructing and Analysing Models	72
8.1	Point-to-Multipoint Communications	72
8.2	Existential Quantification	76
8.3	The Performance of Exhaustive Analysis	76
8.3.1	The Size of the Encoded State Space	76
8.3.2	Execution Times of Exhaustive Reachability Analysis	78
9	Modelling Computer Programs	80
9.1	Data Types	80
9.1.1	Expressive Power	82
9.1.2	Representation	82
9.2	Message Queues	82
9.2.1	Previous Approaches	83
9.2.2	Algebraic Support for Queues	83
9.3	Dynamic Resource Allocation	84
9.3.1	Process Creation	84
9.3.2	Memory Allocation	85
9.4	Procedure Calls	86
9.4.1	Scoping	86
9.4.2	Recursive Procedures	87
9.4.3	Exception Handling	87
9.5	Object-Oriented Constructs	87
9.5.1	Inheritance and Polymorphism	88
10	Conclusion	89
10.1	Contribution of the Thesis	89
10.2	Future Work	90
	Bibliography	91

List of Figures

2.1	Actions in $k = 4$ Steps by $n = 2$ Processes	9
2.2	Atomic Actions in $k = 4$ Steps by $n = 2$ Processes	10
3.1	Automatic Verification of Computer Software	15
4.1	Graphical Representation of a Petri Net	25
5.1	The Union and the Intersection of Ranges	35
7.1	Transition Instance Analysis	52
7.2	An Example Model for Transition Instance Analysis	54
7.3	An Expression $a \cdot a + a \cdot a$ as a Tree and as a Graph	58
7.4	Treating Structured Values in C	61
8.1	Algebraic System Net Model of a Synchronisation Protocol	74
8.2	Graphical Representation of the Model in Figure 8.1	75
8.3	MARIA Model of a Distributed Data Base System	77

List of Tables

7.1	The Interface to the Compiled Expression Evaluator	62
8.1	Encoded State Space Sizes for the Distributed Data Base Model .	77
8.2	State Space Generation Times for the Distributed Data Base Model	78
9.1	Data Types Representable in Different Reachability Analysers . .	81

List of Symbols

$<_{\mathcal{D}}$	total order on \mathcal{D}
\approx_v	term or value compatibility (instance analysis)
A_v	assignment candidates (instance analysis)
$\mathcal{A} = \langle \mathcal{D}^{\mathcal{A}}, \mathcal{F}^{\mathcal{A}} \rangle$	algebra
\mathcal{D}	data type
$\mathcal{D}^{\mathcal{A}}, \mathcal{D}_s^{\mathcal{A}}$	(family of) sorted carriers
$\check{\mathcal{D}}^{\mathcal{A}}, \check{\mathcal{D}}_s^{\mathcal{A}}$	$\mathcal{D}^{\mathcal{A}} \cup \{\varepsilon\}$
ε	undefined symbol
$e_v^{\mathcal{A}}(T)$	evaluation of a term T under the assignment v
\mathcal{F}	function symbols
$\mathcal{F}^{\mathcal{A}}$	operations
$\mathcal{F} \subseteq (\mathcal{T} \times \mathcal{P}) \cup (\mathcal{P} \times \mathcal{T})$	flow relation of a net \mathcal{N}
\mathcal{G}	short-circuit symbols
$i: \mathcal{P} \cup \mathcal{T} \cup \mathcal{F} \rightarrow \mathbf{T}^{\mathbf{S}\mu}$	net inscriptions
\mathcal{M}_{Σ}	the markings of an algebraic system net Σ
$M_0, M \in \mathcal{M}_{\Sigma}$	(initial) marking
$\mathcal{N} = \langle \mathcal{P}, \mathcal{T}, \mathcal{F} \rangle$	net
$o_{\mathcal{D}}: \mathcal{D} \rightarrow \{0, \dots, \mathcal{D} - 1\}$	ordering mapping for data types
$\mathcal{P}, \mathcal{P}_s$	(family of) sorted places
$\Sigma = \langle \mathcal{N}, \mathcal{A}, \mathcal{V}, i \rangle$	algebraic system net
$\mathbf{S} = \langle \mathcal{S}, \mathcal{F}, \mathcal{G} \rangle$	signature
$\mathbf{S}_{\mu} = \langle \mathcal{S}, \mathcal{F}, \mathcal{G}, \mu \rangle$	multi-set signature
\mathcal{S}	sorts
$\mathcal{S}_{\beta} \subseteq \mathcal{S}$	basic sorts
$\mathcal{S}_{\mu} = \{\mu(s) \mid s \in \mathcal{S}_{\beta}\}$	multi-set sorts
$s \in \mathcal{S}$	a sort
$\mathbf{T}^{\mathbf{S}}(\mathcal{V}), \mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$	(family of) sorted \mathbf{S} -terms over \mathcal{V}
$T \in \mathbf{T}^{\mathbf{S}}(\mathcal{V})$	a term
\mathcal{T}	transitions
t_v^+	output effect of a transition $t \in \mathcal{T}$
t_v^-	input effect of a transition $t \in \mathcal{T}$
$\mu: \mathcal{S}_{\beta} \rightarrow \mathcal{S}_{\mu}$	sort mapping
$\mu \in \mathcal{M}(A)$	multi-set over the set A
$\mathcal{V}, \mathcal{V}_s$	(family of) sorted variables
$V^{\mathcal{A}}(\mathcal{V})$	assignments over \mathcal{V}
$\hat{v}, v \in V^{\mathcal{A}}(\mathcal{V})$	input and output assignment of a transition

Part I

Identifying the Problems

Chapter 1

Introduction

The rapid development of telecommunications has considerably increased the use of concurrent and distributed systems. Such systems are much more difficult to manage than stand-alone systems controlled by sequential program code. Concurrent and distributed systems call for new analysis tools, because those used for sequential programs are not very useful in an environment where errors are difficult to reproduce due to the asynchronous nature of the communication between processes.

One solution to this problem is the use of formal models, investigating them in reachability analysers. This technique ensures that the complete state space of the system is checked for certain properties and that complete execution traces to possible errors can be recorded.

The big problem with this approach is, however, the vast state space of all real systems. Therefore, reachability analysis should be applied at an early stage of the development, preferably already when the system is specified. In the specification, there usually are few implementation details that would complicate the analysis and explode the state space. It is also easier and cheaper to correct errors at this stage.

A problem that has not been addressed much is the creation of the model. This is a main obstacle that prevents the introduction of this analysis in the industry in a large scale. Existing analysers often require that a formal model is constructed by hand. This means that the creator must be an expert both in the real system design and in the formalism the analyser uses. Such experts are very hard to find.

One solution is to generate the models automatically, to translate the specification into a formal model. This requires that the specification language has solid semantics and does not leave any room for interpretation. There will always be a need for experts, but their workload can be reduced and the coverability of formal analysis techniques greatly improved by introducing appropriate computer aided tools.

There has been some work in this field [25, 33, 42] connected to the TNSDL programming language [52] used in digital telephone exchanges. A major difficulty was that the reachability analyser [62] did not support the data types used in the specification language. Thus the decision was made to design a new analyser that supports a large class of data types [41] and lets the modeller or translator

writer concentrate on other things than the representation of data.

This work is about the design of data types for the formalism used in the MARIA analyser and about the implementation of the analyser. The analyser was designed to meet the demands of industrial systems. It is important to notice that while the data type system has been designed in such a way that it satisfies the *practical needs* of specification and programming languages, the modelling language of the analyser is *completely formal*.

The formalism the analyser applies is a variant of Algebraic System Nets [36, 37, 51] with a fixed set of algebraic operations. Algebraic System Nets, inspired by many-sorted nets [3] and coloured nets [31], are a high-level enhancement of ordinary Petri Nets [49, 50].

Throughout the work we have strictly followed the requirements of reachability analysis, so that all models that can be entered to our tool can be analysed in a formal way. The whole reachability analyser has been designed with the following requirements in mind:

1. satisfaction of practical needs, ease of use
 - structured data types
 - aggregate operations on multi-sets
2. reachability analysis
 - strictly formal semantics
 - very tight representation of states
 - theoretical possibility to unfold all models to low-level Petri Nets
3. modular design; replaceable modules for
 - expression evaluation and handling
 - reachability analysis and model checking
 - reachability graph management

When constructing a formal model of a system, one should pay attention not only to correctness but also to possible difficulties in reachability analysis. The expressive power of the formalism is not merely important for easy modelling but also for efficient analysis. If the formalism contains high-level data types like queues and stacks, the set of reachable states can be reduced considerably.

The definition of Algebraic System Nets given in [37] was quite useful for this work, but some additions had to be made in order to achieve better expressiveness and efficiency in the analysis. For instance, our implementation allows multi-set terms with non-constant multiplicity and a special case of multi-set valued variables. Algebraic System Nets define a fairly generic framework for describing computations; this work defines a set of concrete data types and algebraic operations that should be adequate for the succinct description and efficient analysis of practical concurrent systems.

Designing data types for a formal analyser is quite different from designing data types for a programming language. The data types must not only be compatible with the formalism but also with the analysis techniques. Thus it was impossible to allow general pointers in our formalism, as we shall see in Chapter 5. Unbounded data types like lists and trees were also omitted. However, after careful examination it was possible to include most of the data types used in common programming languages. It was also possible to impose a total order on all data types—a useful feature in the analysis and a requirement for some powerful algebraic operations.

1.1 Reachability Analysers in the Past and Present

Formal analysis tools based on exhaustive state space exploration, or reachability analysers, have greatly developed during the past two decades. There are numerous research groups both at universities and in commercial companies, and many tools have been developed for internal use, or just to see whether a theoretical idea might work in practice, often analysing theoretical models that do not directly have any roots in the real world.

The work on automated protocol validation and tools for Petri Nets began in the 1970's. In the beginning, when the memory capacity of computers was severely limited, reachability analysis could only be applied to rather simple systems. In the early 1980's, the most advanced analysers could handle systems of up to tens of thousands of reachable states, and model checking—checking whether there are execution paths that violate a property expressed in temporal or modal logic—was considered impractical by some researchers. [65]

In Finland, one of the first research projects on computer tools for reachability analysis started in the summer of 1980. Financed by the then Posts and Telecommunications of Finland, the three-year project resulted in a theorem prover for modal logic [38]. In 1984, the research continued in a joint four-year project of the Computer Technology Laboratory of Technical Research Centre of Finland (VTT) and the Digital Systems Laboratory of Helsinki University of Technology. Financed by the then Technology Development Centre of Finland (TEKES) and four industrial partners, the Rimst project (Rinnakkaisjärjestelmien määrittelyn ja suunnittelun tukijärjestelmä, or Support System for Specification and Design of Concurrent Systems) identified a number of practical and theoretical problems, some of which could be addressed during the project. When the project ended in 1988, the Digital Systems Laboratory continued to explore Petri Nets, while the group at VTT eventually developed a tool set for labelled transition systems [34].

At Helsinki University of Technology, the joint project resulted in a set of analysis tools for Predicate/Transition Nets [14] called PRENA [35]. The worst limitation of the tool was that it only could handle some thousands of reachable states. By that time, there were several reachability analysers in existence. A survey from 1988 [39] lists 22 tools for high-level Petri Nets, 13 of which perform reachability analysis. An earlier survey from 1985 [11] that includes tools for low-level nets lists 26 tools, most of which have been written in nonportable languages

for proprietary systems. The situation seems to have stabilised a little: a survey from 1998 [55] focuses on ten tools for high-level nets.

The initial version of PROD, the successor of PRENA, was developed in 1989–1991 mostly as a student project. Capable of exhaustively analysing systems with millions of reachable states and performing model checking while generating the set of reachable states, the tool has been a flagship of the Digital Systems Laboratory, later Laboratory for Theoretical Computer Science, and it has been extended with advanced algorithms, such as a model checker for the branching time temporal logic [22] and the stubborn set method [61].

Difficulties in modelling some practical systems presented the need for a reachability analyser that supports a more powerful modelling formalism. One of the main goals of the three-year MARIA project, which started in 1998, was to develop a reachability analyser that supports a highly expressive data type system. This goal has been achieved: our tool is capable of performing brute-force exhaustive reachability analysis both interactively and in batch mode. Still, much work remains to be done, and many of the advanced algorithms of PROD are being implemented in MARIA. There are also some promising reduction methods that have not been implemented in either analyser yet, such as the symmetry reduction method [32].

There are many institutions with a long tradition in reachability analysis. At Bell Labs, reachability analysers have been developed and used since the early 1980's. Their current analyser SPIN [24] is designed for analysing computer protocols, and it is one of the few analysers that have successfully been applied to industrial-size systems. The tool is based on a process-oriented modelling language PROMELA whose syntax resembles some programming languages. The somewhat informal application-oriented approach of SPIN has turned out to suit extremely well to modelling protocols.

SPIN supports probabilistic verification, representing the set of reachable states with a very large hash table, a bit vector indexed by hash values. The larger the table and the better the hash function that converts system states to indices of the hash table, the smaller is the probability that an unexplored state is mistaken for an already explored one. When using this probabilistic method, one cannot be completely sure that the whole state space has been explored. Nevertheless, it is hard to match SPIN in performance, especially with a tool meant for analysing all kinds of concurrent systems and not only protocols.

SDL, the CCITT Specification and Description Language [30], which is mainly used by the telecommunications industry for specifying protocols, has raised some interest both in the academic world and among commercial tool vendors. The verification tool by Telelogic [10] is probably the most capable reachability analyser for SDL, when it comes to the extent of supported language constructs. However, the tool does not appear to be able to perform model checking of properties expressed in temporal logic, and it generates the reachability graph in the system memory, which makes the exhaustive analysis option useless for nontrivial specifications. In academic projects, small subsets of SDL have been translated to the internal formalisms of some reachability analysers, such as PEP [19] and PROD [62].

1.2 Related Work

Not all tool authors aim for the ability to efficiently perform exhaustive reachability analysis and model checking. For teaching purposes and for running simulations, a graphical user interface and an inscription language with a lot of expressive power seem to be more important than a solid theoretical background.

One tool, DESIGN/CPN [43], is based on a formalism known as Coloured Petri Nets [31]. While the tool is capable of performing exhaustive reachability analysis, the high expressive power of its inscription language—including the ability to define data types of an unbounded domain—makes it very difficult to unfold the net to a lower level, which is the domain of many advanced analysis methods. Models constructed in our formalism can always be unfolded in principle; see Section 7.1.2.

The efforts in modelling high-level languages using Petri Nets have shown that the underlying net formalism must have enough expressive power to facilitate a straightforward transformation. Otherwise the transformation is awkward [33, 42], or the domain of the source language has to be severely restricted [20], or both, as we illustrate in Chapter 9. However, the introduction of new constructs must be carefully considered in order not to sacrifice tractability for expressiveness.

1.3 Outline

We begin with some general observations on the properties of typical computer programs and suggesting how they should be exploited in formal analysis. In Chapter 3 we list some of the techniques that can be applied in order to detect errors in different stages of software development.

The second part of this work, along with the third part, is the main contribution of the thesis. We introduce a class of many-sorted algebras and define Algebraic System Nets in Chapter 4. Chapter 5 gives an inductive definition of the concrete data types implemented in our reachability analyser [41] for Algebraic System Nets, and Chapter 6 defines the algebraic operations.

In the third and last part of the work, we discuss some implementation details, such as finding all enabled actions in a state, evaluating algebraic terms efficiently, and representing the states and enabled actions of a system as very short bit vectors. Finally, Chapters 8 and 9 motivate some of the design choices of our analyser by showing how certain common constructs, which are difficult to represent in other formalisms, can be translated to our class of nets in a straightforward way.

Chapter 2

Verification of Concurrent Programs

Many concurrent and distributed systems include components implemented in computer software.¹ Due to the complex nature of such systems, it is difficult to convince oneself of their proper operation under all circumstances just by looking at the program code. Concurrency related errors that manage to go undetected in the manufacturer's tests can be expensive to correct. The manufacturer would rather experience a system crash in the laboratory and not at the customer's premises.

2.1 Classifying Programs

In the software industry, there is a famous law referred to as the 20 %–80 % rule or the 10 %–90 % rule. Usually it refers to the size of a program and to its execution time: when a program is executed, most of the time is spent executing statements in a small fraction of the program text. This law could also be used to classify concurrent programs. Typically, only a small part of a concurrent program deals with concurrency, and assuming that the errors in the program are evenly distributed and that formal methods are not being applied, most of the time devoted to debugging will be spent trying to figure out why the concurrent algorithm fails under some circumstances.

2.1.1 Sequential Behaviour

Errors in the sequential parts of a computer program are relatively easy to locate. Modern compilers perform quite a lot of static analysis and can issue warnings for things like unused or uninitialised variables, for enumeration constants not handled in a switch statement, and so on.

Static checking can only detect a small class of errors. In practice, the majority of other errors in sequential code can be found by testing. Since sequential programs typically are fully deterministic, i.e. they always produce the same output given the same input, it is easy to reproduce erratic situations and to find their

¹Verifying the correct operation of hardware circuits is beyond the scope of this work.

cause by stepping through the code in a debugger. This approach, although it is far from formal, works astonishingly well for sequential program code.

2.1.2 Concurrent Behaviour

Constructing or understanding event or message driven programs requires a different way of thinking than traditional algorithm driven programs. The event driven approach is common in reactive or embedded systems, in graphical user interfaces and in applications working over computer networks. As the input events or messages are typically generated by entities running concurrently with the system, the system as a whole tends to behave in a nondeterministic and irreproducible way.

Exhaustive analysis of the concurrent behaviour of a system calls for the use of formal methods. Tests or simulations can only prove the presence of a bug (by accidentally finding one), but only exhaustive formal analysis can prove the absence of bugs in an abstract model of the implemented system.²

2.2 Making Abstractions

Although it is possible to make a formal model of any program written for a finite number of finite-memory computers, it is often undesirable to do so, since the set of reachable global states of the system can be excessively large.

Large systems can be formally analysed by omitting details that are considered irrelevant or uninteresting. There are many different ways to make abstractions; some alter a formal model, others affect its interpretation. In this section, we shall give two examples of making abstractions in a model.

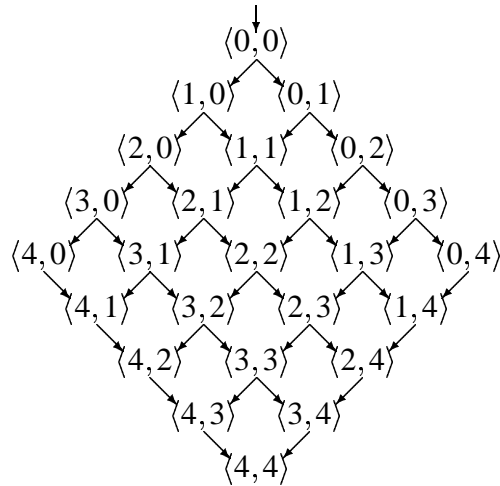
2.2.1 Atomising Sequences of Actions

Consider a system consisting of n concurrently executing processes. If each process performs k actions before communicating with other processes, the system will have $(k + 1)^n$ possible global states and $nk(k + 1)^{n-1}$ possible transitions between them. Figure 2.1 illustrates the reachability graph of such a system. The example is from the introductory part of Antti Valmari's dissertation [59, Section 3.5.1], which also contains proofs for the above numbers.

The k actions of each process could represent purely sequential behaviour, which can be abstracted away, efficiently letting $k = 1$. Still, the size of the thus abstracted state space grows exponentially with n . If we make the assumption that the processes execute synchronously, we have $n = 1$ and get a easily manageable number of states.

Since errors in sequential program code can be treated pretty well by using semi-formal techniques such as testing, it is reasonable to assume that there are no

²Also formal analysis can produce wrong results if the model does not adequately represent the system and its environment.

Figure 2.1: Actions in $k = 4$ Steps by $n = 2$ Processes

errors in the sequential parts when constructing a model for verifying the concurrent behaviour of a system. Thus, the parts internal to processes can be abstracted away from the model without losing any suspicious behaviour.

In our example system of n processes performing k actions, the actions are internal to each process. Since the processes do not communicate with each other until each of them have completed the internal actions, there is no way the internal state of one process could affect the behaviour of other processes.

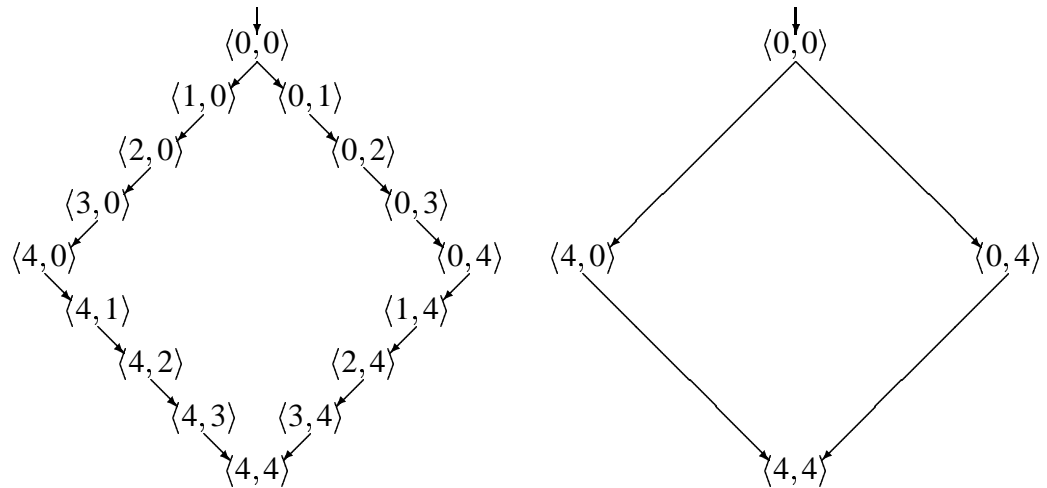
In the lattice-shaped reachability graph presented in Figure 2.1, the vertices in the middle represent intermediate states where each process has performed some but not all of its actions. If the actions are treated atomically—that is, once a process starts performing a sequence of internal actions, it will complete the whole sequence before other processes do anything—we will get rid of those intermediate states.

Figure 2.2 illustrates the effect of atomising sequences of local events. It is straightforward to see that instead of $(k + 1)^n$ states and $nk(k + 1)^{n-1}$ transitions, the reachability graph now only has $(n(k - 1) + 2)2^{n-1}$ states and $nk2^{n-1}$ transitions. Furthermore, if the sequences of events are collapsed to single events as shown in the right part of the figure, we will end up with 2^n states and $n2^{n-1}$ transitions.

This simple example shows that the reachability graph of a concurrent system can be vastly reduced by optimising the model. Further savings can be achieved during the reachability graph generation with advanced analysis techniques such as partial order reductions and symmetry reductions, which are summarised in [60].

2.2.2 Introducing Nondeterminism

It may seem that all concurrent systems can be efficiently modelled by converting complex local computations to single events, thus pruning most intermediate states. Unfortunately this turns out not to be the case. Some analysis techniques

Figure 2.2: Atomic Actions in $k = 4$ Steps by $n = 2$ Processes

become utterly inefficient when the source and target state of a transition differ very much or when extremely many transitions are possible in one state. Introducing a few intermediate states could be of some help, but there is a better solution.

The sequential part of a concurrent system usually constrains and limits the behaviour of the system. For example, in a distributed game, the state of the game field and some decision algorithms will determine the type of messages sent to different counterparts. When we are only interested in generic properties like the absence of deadlocks in the communications protocol of the game, we can abstract away all information regarding the game field (also from the messages sent by different processes) and model the outcome of the decision algorithm by a non-deterministic choice of all possible outcomes, e.g. “game over” or “your turn.”³

Omitting the typically large data structures needed by the sequential parts of a system from the abstract verification model will not merely reduce the amount of memory needed for storing the global state of the system; also the computational complexity involved with the reproduction of the exact behaviour of the sequential part will be reduced to a trivial linear-time operation of making a nondeterministic choice between all possible outcomes of the computation. Often also the outcome can be presented at a higher abstraction level, which can drastically reduce the domain.

The use of libraries characterises modern programming. It is difficult to imagine a large object-oriented program that does not make any use of built-in libraries for basic data structures, such as lists and search trees. Sometimes libraries are part of the run-time environment and completely transparent for the user. For instance, the virtual machine of the popular interpreted language Perl [63] originally designed for text processing has an instruction that matches character strings by generalised regular expressions [53, p. 324]. Including all these complex al-

³The original decision algorithm would probably have several outcomes equivalent to “your turn” in the abstracted version. As the state of the game field has been abstracted away, there is no need to distinguish between “I pass” and “I move the piece from place x to place y .”

gorithms in the abstract model of a system would be like reinventing the wheel and certainly not worth the effort. It is rather unlikely that the library routines of a widely used language implementation contain errors.

Hidden Errors

There is no free lunch, not even in the world of abstractions. Consider a ring network where messages are relayed from node to node. If the communications protocol has been built so that a node may only send a message to another once it has received a message either for itself or for relaying, it can happen that a node has to wait for a long time before it can send a message. If the abstract model of the nodes is totally nondeterministic, so that any node can decide to send a message to another node at any time, the problem will fail to arise in the formal analysis.

Also abstracting chunks of sequential code can hide severe errors. If the original piece of code under some circumstances can break the bounds of an array or be trapped in an infinite loop, such errors will be absent in a model that replaces the chunk with a nondeterministic choice.

False Alarms

Abstractions can not only hide errors; they can also cause false alarms. For instance, consider the distributed game described earlier. The communications protocol may have been specified so that it is an error to begin the game with the message “game over.” Replacing the decision algorithms with nondeterministic choices can make it possible to start the protocol with that message. If the model has some kind of a guard against this, the analysis will yield a false alarm.

False alarms may seem harmless in comparison to hidden errors, but there is a problem involving the human element. If the person who interprets the results of the analysis receives an extensive list of error traces, he will probably check the realisability of some of them and ignore all further error messages of that type. Realisable errors can then also go unnoticed.

Chapter 3

Computer Tools for Analysis

Applying formal methods in software development may seem straightforward. All sequential behaviour will be modelled with nondeterminism, and only the fraction of the system expressing concurrent behaviour must be translated to an abstract model. But the reality can sometimes be quite different. If the project manager wants to have a 50,000-line program verified by tomorrow or even by next week, manually constructing a verification model is not an option, especially if the work has to be done by someone who is not very familiar with formal methods or with the formalism applied by the analysis tool.

It is sad but true that in many software development projects, quality issues and formal methods are typically forgotten until the project encounters problems too severe to be solved in a reasonable time period of hacking and debugging. Designing the system incrementally, verifying and testing each refinement step, will help to detect errors earlier. For the computer tools described here, however, it does not make any difference whether they are used from the very beginning of a project or as a last resort to patch a sinking ship sailing on a stormy ocean.

3.1 Incomplete Methods

Practical systems tend to be so complex that exploring the complete state space, checking that the system performs correctly on all possible input sequences, is out of the question. Many errors can be found with less ambitious methods, some of which will be described here.

3.1.1 Static Analysis

Some errors can be discovered by relatively simple analysis. In the early days of computing when memory was scarce, compiler front-ends were rather simple, and separate programs were developed for detecting simple mistakes such as using the assignment operator where an equality comparison is likely to be intended.

Modern compilers perform all kinds of optimisations, which require rather thorough static analysis of the input. Many sanity checks are performed as a side effect. It is common to check for unreachable program statements or for expres-

sions whose value is not used, and often warnings about them indicate typing mistakes or even an error in the program logic.

Many compilers have some very strict static analysis options that are disabled by default, since they would cause numerous warnings for existing programs or even for the built-in libraries. Enabling such options only makes sense if the program does not need too big modifications in order to pass the checks. A good example is the check for uninitialised member variables by a C++ [27] compiler, which requires that all constructors make use of the member initialisation list, a less known feature of the language. Normally it makes sense to filter out certain warning messages reported for library interfaces, so that only relevant, correctable mistakes will be pointed out.

3.1.2 Instrumenting Program Code

Errors in large programs tend to have global effects. For instance, if an algorithm corrupts some data structures, the corruption may cause another, correctly implemented algorithm to fail. It depends on the underlying safety net, e.g. on the granularity of memory protection, how far the effects will propagate before the system crashes.

A tight safety net that would e.g. check the validity of arguments and return values of all procedure invocations will detect errors closer to their cause. Also in this case there is no free lunch: a program *instrumented* with all kinds of sanity checks will run slower, and if the program is free of errors, all the checks are redundant and unnecessary. Usually programs are instrumented during the development and testing phase, and the checks are omitted from the published version if the performance penalty is an issue.

Assertions

The C programming language [28] defines a macro for making *assertions*, for ensuring that a condition holds whenever the macro is evaluated. The macro is supplied with one argument, a truth-valued expression. If the expression evaluates to false, an error message will be displayed and the execution of the program will be aborted. Assertions can be disabled by defining a preprocessor symbol when compiling the program. Therefore, the same source code can be used for producing a fast, optimised executable as well as an instrumented executable for testing and debugging.

Detecting Resource Leakages

A program that allocates a resource but does not ever give it up is said to have a *resource leakage*. Probably the most common type of resource leakages is related to dynamic memory allocation. In programming languages whose run-time environments do not automatically release unused dynamically allocated memory in a process called *garbage collection*, it is the programmer's responsibility to explicitly deallocate memory areas that are no longer needed.

The need for explicit memory deallocation causes two major problems. First, programs that constantly allocate memory but do not deallocate it will eventually run out of memory. Second, if a program is too eager about deallocating memory, it can end up with dangling pointers pointing to deallocated memory. Accessing deallocated memory through these pointers or trying to deallocate a memory block more than once will probably corrupt data structures and may cause a crash somewhere else in the program code.

These problems can be addressed with programs called memory access debuggers. They work either by replacing the library routines for memory allocation and deallocation [48], by instrumenting the compiled code [9] or by interfacing with the compiler and by instrumenting the program at compile time, a technique widely applied both by free [15, 64] and by commercial tools.

Memory access debuggers have severe limitations. The instrumented program typically executes an order of magnitude slower than the original, and it will also require much more memory. This is due to the fact that memory access debuggers typically do not deallocate the memory deallocated by the instrumented program. By doing so, they can detect accesses to deallocated memory blocks. Also, the instrumented program code will allocate bigger memory blocks than the program actually asks for, so that it will be possible to detect accesses outside the bounds of the allocated memory area.

3.1.3 Regression Testing

One way to make software development resemble a controlled, deterministic process rather than the efforts of a panicking fire brigade is systematic testing. Whenever new features are added to a piece of software, a set of test cases are defined. The test cases, consisting of input sequences and expected output sequences, are stored in a version control system. Before accepting modifications to the program, it must pass all tests recorded so far. If it does not, the program has to be corrected or the test case must be updated.

The key problem with regression testing is the need for manual intervention. It is difficult to design tests that cover all of the program code. Moreover, when a test fails, one must ensure that it is not the test case that violates the specification, which often is not carved in stone but lives with the development cycle. In any case, it must be kept in mind that testing can only prove the presence of errors, not the absence of them.

3.2 Formal Methods

Hardly anyone believes that the correct operation of a complex system could be formally verified simply by pushing a button. Expertise in formal methods will always be needed to be able to specify the properties the system should fulfill, to interpret the results when a complicated error is detected and to simplify the model of the system by means of abstractions. However, there is no reason why

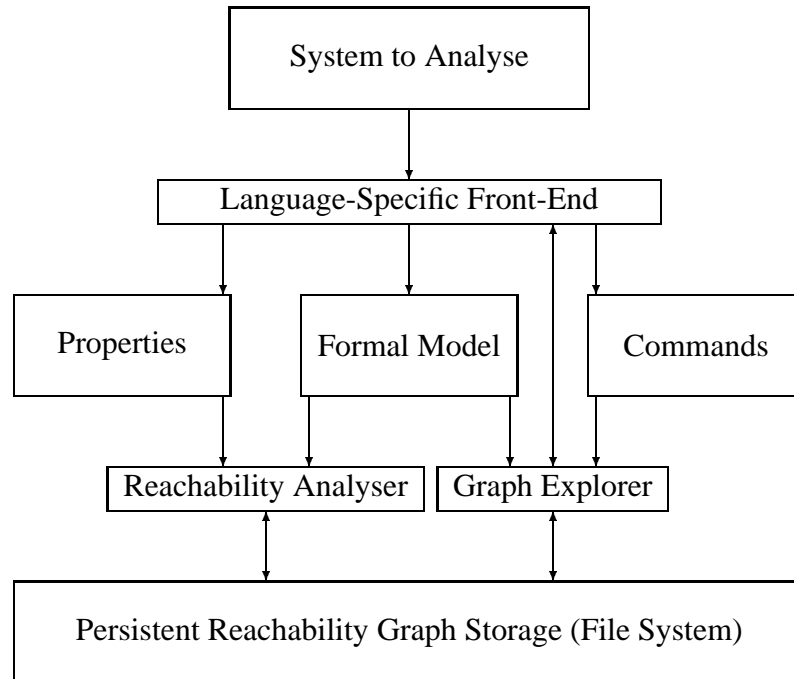


Figure 3.1: Automatic Verification of Computer Software

verification could not be made to appear as a push-button technology for system designers and engineers. It is enough to have experts behind the scenes.

Figure 3.1 illustrates our vision of applying formal methods to software verification. There are several reasons why computer aided verification should be automated like this instead of manually constructing models compatible to verification. First, it is not reasonable to assume that everyone can learn formal methods to the necessary extent. Only a few people would have the necessary skills to construct the model. Second, modelling big systems involves much work, and humans make mistakes when performing mechanical tasks. Also, if the system is updated frequently, it would be hard to keep the model up to date.

3.2.1 Constructing the Model

Since most designers are not familiar with other formalisms than the one they are directly working with, the verification model has to be hidden from them by means of a *language-specific front-end*. The front-end will translate the implementation expressed in an application-specific high-level language to an analysable formal model, making reasonable abstractions. The model will be checked for the absence of deadlocks and for properties that have been entered in a database by an expert. Counterexamples, that is, execution paths leading to a state violating a desired property, will be translated back to the original formalism. In this way, the average user does not need to know the details of the underlying machinery.

Automatic model construction has many prerequisites. The system has to be implemented in a reasonably high-level language that has a well-defined se-

antics. If the system description cannot be represented with a single abstract syntax tree, it is difficult to construct a formal model of it. In practice, a system can be implemented using several programming languages and some electronic circuits. Models for low-level subsystems have to be constructed manually. Also the abstraction rules for omitting certain things from the model have to be specified by an expert. The databases of desired properties and abstraction rules may need to be updated during the development process if the specification changes.

Modelling Data

Incorporating big amounts of data in models intended for verification is generally a bad idea, since it raises the memory requirements for storing the global state of the model and often splits logically equivalent states to a number of states, multiplying the number of nodes in the reachability graph. Also, complex data structures are typically maintained by sequential code, which should be omitted from the model for reasons given in Chapter 2.

Nevertheless, all models need some data, and if a system is simple enough, it can be formally analysed without omitting anything. Sometimes it is nice to experiment with a model, to see which parts need to be abstracted in order to be able to analyse the behaviour. This kind of prototyping and experimenting requires that the data types and algebraic expressions used in the system can be automatically translated to corresponding structures in the model.

If a system implementation uses a construct that cannot be directly expressed in the modelling language, constructing an automatic transformation can be a major challenge, especially when attention is paid to the performance of the generated model. Data types are very problematic in this aspect. It is cumbersome to represent arrays or first-in-first-out buffers in a formalism that is based on tuples of integers or enumerated constants. Doing so may be of academic interest [20, 25], but if the intention is to accomplish some real work, applying an inefficient transformation that only works in special cases will clearly be out of question.

A modelling language equipped with powerful data types such as structures, unions, arrays and variable-length buffers is an easy target formalism for compilers of programming languages. If the automatically constructed model generates too big a state space, simplifications and abstractions can be specified in the source formalism or in the abstraction rules. This can be done by someone familiar with the system but not necessarily knowing the modelling formalism.

Modelling the Environment

In order to analyse a program, it has to be given some input. In distributed systems, the input typically consists of events generated by the environment. Usually the environment does not act randomly, but it follows some disciplines. For instance, if one wants to place a telephone call, he will first dial the number and will not start talking until the call is answered.

The environment has to be part of the verification model, for various reasons. Connecting the model to a totally nondeterministically acting environment, which

can synchronise with all kinds of actions at all times, could result in an unmanageable number of states and in a large number of errors that are impossible in the real system. Also, a nondeterministic environment may drive the system to erroneous states that are impossible in practice, or let the system continue from a situation that would be a deadlock if there were some constraints for the environment.

3.2.2 Analysing the Model

Merely constructing a formal model of a system does not solve anything. The model has to be analysed to see whether it (and the system it represents) fulfills some desired properties. Analysis can be carried out at different levels; here we will present some possibilities.

Reachability Analysis

Exhaustive reachability analysis, considering each possible action in each reachable state of the model, is the most complete analysis method. Its major drawback is that for many practical systems, the set of reachable states tends to be so large that the state space cannot be efficiently stored in a computer memory. There are several techniques that address this *state space explosion* problem; see [60] for an introduction.

Model Checking

Only a limited set of properties, such as the presence or absence of deadlock states, can be found out from the reachability graph generated by exhaustive reachability analysis. Model checking can prove or find violating execution paths for properties typically expressed in temporal logic. If a property does not hold, a counterexample will usually be found after generating only a part of the reachability graph.

Simulation

Like software development, also modelling can be an incremental, experimental process. Simulation is a very useful tool that lets one to experiment with a model, to make sure that it behaves in the intended way. An incorrectly working model will often produce an infinite state space. Performing exhaustive reachability analysis on such a model would be a waste of time; often hours or even days can pass before the analyser will run out of memory.

Simulation is also very useful for illustration and teaching purposes. In simulation, only a small subset of possible actions will be carried out. Simulations typically work on one trace, always performing an action on the state generated by the previously performed action. The next action can be chosen either randomly among the set of possible actions, or it can be chosen by the user.

A somewhat different simulation approach has been implemented in [41]. Instead of letting the user to choose actions, the analyser will perform all possible actions that are possible in the state picked up by the user. This kind of simulation will not produce a single trace, but a (partial) reachability graph, which can

be explored at all times. It is also possible to evaluate temporal logic formulae in a state. Only if the model checker does not find a counterexample for the formula, the complete reachability graph will have to be generated, unless the query is interrupted. Interactive on-demand reachability graph generation appears to combine the advantages of simulation with the advantages of reachability analysis at a relatively small price.

Part II

Designing the Formalism

Chapter 4

Algebraic System Nets

Our formalism for modelling and analysing concurrent systems is the class of Algebraic System Nets [36, 37, 51], which we will describe in this chapter. Later parts of this work will refine and restrict the formalism, as used by our implementation [41] of a reachability analyser.

4.1 Signatures and Algebras

Algebraic System Nets, as their name suggests, are based on algebras. The algebras we represent here contain two extensions: error checking and short-circuit evaluation. Our notion of errors, inspired by [32], bears similarities with exception conditions [18] and abstract errors [17]. To keep the notation simple, we do not distinguish between different kinds of errors.

The motivation behind the second extension of our algebras, short-circuit evaluation of algebraic terms, is to narrow the gap between algebras and optimised computer implementations of expression evaluators.

Languages have two major properties: syntax (appearance) and semantics (interpretation). For algebraic systems, the syntax is defined in a signature, which contains symbols and terms but does not define any interpretation for them. The semantics is given in the algebra, which consists of the signature and of functions corresponding to the operation symbols given in the signature.

Our definition of algebras makes use of a concept of families, collections of sets. We assume that the reader has a knowledge in some mathematical preliminaries such as sets.

Definition 4.1 (Families) *For a set I ,*

$$A = \bigcup_{i \in I} A_i$$

is a family of sets if A_i is a set for each $i \in I$. Furthermore, a family A is pairwise disjoint if for all $i, j \in I$, $i \neq j \Rightarrow A_i \cap A_j = \emptyset$.

4.1.1 Signatures, Variables and Terms

Definition 4.2 (Signatures) A signature

$$\mathbf{S} = \langle \mathcal{S}, \mathcal{F}, \mathcal{G} \rangle$$

consists of

1. a non-empty set \mathcal{S} of sort names (sorts)
2. a pairwise disjoint family $\mathcal{F} = \bigcup_{\sigma \in \mathcal{S}^*, s \in \mathcal{S}} \mathcal{F}_{\sigma, s}$ of functional operation symbols
3. a pairwise disjoint family $\mathcal{G} = \bigcup_{s', s \in \mathcal{S}} \mathcal{G}_{s', s^n}$ of short-circuit operation symbols where s^n stands for $\underbrace{s, \dots, s}_n$; $\mathcal{F} \cap \mathcal{G} = \emptyset$.

An operation symbol $f \in \mathcal{F}_{s_1 \dots s_n, s}$ stands for a functional operation from the domain sorts s_1, \dots, s_n to the range sort s of f . The set $\mathcal{F}_{\lambda, s}$ is called the set of **S-constant symbols of sort s** , where λ denotes the empty sequence.¹

An operation symbol $g \in \mathcal{G}_{s', s^n}$ stands for an operation from s' to an operation from s^n to s , where s' is the *selection sort* and s the *range sort* of g .

Definition 4.3 (Variables) A pairwise disjoint family

$$\mathcal{V} = \bigcup_{s \in \mathcal{S}} \mathcal{V}_s$$

of symbols is called a family of **S-variables**.

Using variables and the two kinds of operation symbols, we can build **S-terms**, sequences of symbols, according to the following definition, which defines a kind of grammar. Algebraic terms can be viewed as strings of symbols; their interpretation is defined separately.

Definition 4.4 (Terms) The set $\mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$ of **S-terms** of sort $s \in \mathcal{S}$ over \mathcal{V} is the minimal set defined inductively by the following rules.

1. $\mathcal{V}_s \subseteq \mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$.
2. For $n \geq 0$, if $f \in \mathcal{F}_{s_1 \dots s_n, s}$ and $T_i \in \mathbf{T}_{s_i}^{\mathbf{S}}(\mathcal{V})$ for $1 \leq i \leq n$, then $f(T_1, \dots, T_n) \in \mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$.
3. For $n \geq 0$, if $g \in \mathcal{G}_{s', s^n}$ and $T' \in \mathbf{T}_{s'}^{\mathbf{S}}(\mathcal{V})$ and $T_i \in \mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$ for $1 \leq i \leq n$, then $g(T', T_1, \dots, T_n) \in \mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$.

The set $\mathbf{T}_s^{\mathbf{S}}(\emptyset)$ is the set of **S-ground terms** of sort s .

¹By our convention, sequence indices are written in ascending order. When the index of the last element of a sequence is smaller than the index of the first element, the sequence is empty. For instance, if $n < 1$, the sequence $s_1 \dots s_n$ equals λ , the empty sequence.

4.1.2 Algebras, Assignments and Evaluations

An algebra corresponding to a signature gives an interpretation for the sorts and operations of the signature. It assigns each sort a carrier and each functional operation symbol a function. Chapter 5 defines one possible family of carriers (called data types), and Chapter 6 defines a family of operations on them. This chapter proceeds in a more abstract level, independent of the actual carriers and functions.

Definition 4.5 (Short Circuit Error Algebras) *Let $\mathbf{S} = \langle \mathcal{S}, \mathcal{F}, \mathcal{G} \rangle$ be a signature. An \mathbf{S} -short circuit error algebra, or \mathbf{S} -algebra*

$$\mathcal{A} = \langle \mathcal{D}^{\mathcal{A}}, \mathcal{F}^{\mathcal{A}} \rangle$$

consists of:

1. a pairwise disjoint family $\mathcal{D}^{\mathcal{A}} = \bigcup_{s \in \mathcal{S}} \mathcal{D}_s^{\mathcal{A}}$ of non-empty carriers
2. the undefined symbol $\varepsilon \notin \mathcal{D}^{\mathcal{A}}$
3. a pairwise disjoint family $\check{\mathcal{D}}^{\mathcal{A}} = \bigcup_{s \in \mathcal{S}} \check{\mathcal{D}}_s^{\mathcal{A}}$ of augmented carriers; for all $s \in \mathcal{S}$, $\check{\mathcal{D}}_s^{\mathcal{A}} = \mathcal{D}_s^{\mathcal{A}} \cup \{\varepsilon\}$
4. a pairwise disjoint family $\mathcal{F}^{\mathcal{A}} = \bigcup_{f \in \mathcal{F}} f^{\mathcal{A}}$ of functional operations; for all $f \in \mathcal{F}_{s_1 \dots s_n, s}$ with $n \geq 0$, $f^{\mathcal{A}}$ is a mapping

$$f^{\mathcal{A}} : \check{\mathcal{D}}_{s_1}^{\mathcal{A}} \times \dots \times \check{\mathcal{D}}_{s_n}^{\mathcal{A}} \rightarrow \check{\mathcal{D}}_s^{\mathcal{A}}$$

such that the image of the subset

$$(\check{\mathcal{D}}_{s_1}^{\mathcal{A}} \times \dots \times \check{\mathcal{D}}_{s_n}^{\mathcal{A}}) \setminus (\mathcal{D}_{s_1}^{\mathcal{A}} \times \dots \times \mathcal{D}_{s_n}^{\mathcal{A}})$$

equals $\{\varepsilon\}$; that is, whenever an argument equals ε , so will also the result

5. bijective mappings

$$o_{\mathcal{D}_{s'}^{\mathcal{A}}} : \mathcal{D}_{s'}^{\mathcal{A}} \rightarrow \{0, \dots, |\mathcal{D}_{s'}^{\mathcal{A}}| - 1\}$$

for each short-circuit term $g \in \mathcal{G}_{s', s^n}$, where the selection sort s' has a finite carrier $\mathcal{D}_{s'}^{\mathcal{A}}$ with $|\mathcal{D}_{s'}^{\mathcal{A}}| = n$.

Definition 4.6 (Assignments) *Let \mathcal{A} be an algebra. The set*

$$V^{\mathcal{A}}(\mathcal{V}) = \left\{ v \left\| v : \bigcup_{s \in \mathcal{S}} (\mathcal{V}_s \rightarrow \check{\mathcal{D}}_s^{\mathcal{A}}) \right. \right\}$$

is the set of assignments to the variables of the family \mathcal{V} .

Note that we allow also *undefined variables*, which are assigned the undefined value ε . Given an assignment, \mathbf{S} -terms can be evaluated as follows:

Definition 4.7 (Evaluations of Terms) An assignment $v \in V^{\mathcal{A}}(\mathcal{V})$ is extended to the corresponding evaluation of terms

$$e_v^{\mathcal{A}} : \bigcup_{s \in \mathcal{S}} \left(\mathbf{T}_s^{\mathcal{S}}(\mathcal{V}) \rightarrow \check{\mathcal{D}}_s^{\mathcal{A}} \right)$$

in the following inductive definition for each $T \in \mathbf{T}_s^{\mathcal{S}}(\mathcal{V})$:

1. If $T \in \mathcal{V}_s$, then

$$e_v^{\mathcal{A}}(T) = v(T).$$

2. If $T = f(T_1, \dots, T_n)$, $f \in \mathcal{F}_{s_1 \dots s_n, s}$ and $T_i \in \mathbf{T}_{s_i}^{\mathcal{S}}(\mathcal{V})$ for $1 \leq i \leq n \geq 0$, then

$$e_v^{\mathcal{A}}(T) = f^{\mathcal{A}}(e_v^{\mathcal{A}}(T_1), \dots, e_v^{\mathcal{A}}(T_n)).$$

3. If $T = g(T', T_1, \dots, T_n)$, $g \in \mathcal{G}_{s', s^n}$, $T' \in \mathbf{T}_{s'}^{\mathcal{S}}(\mathcal{V})$ and $T_i \in \mathbf{T}_{s_i}^{\mathcal{S}}(\mathcal{V})$ for $1 \leq i \leq n \geq 1$, then

$$e_v^{\mathcal{A}}(T) = \begin{cases} \varepsilon & \text{if } e_v^{\mathcal{A}}(T') = \varepsilon \\ e_v^{\mathcal{A}}(T_{k+1}) & \text{if } \exists k \in \{0, \dots, n-1\} : k = o_{\mathcal{D}_{s'}^{\mathcal{A}}}(e_v^{\mathcal{A}}(T')) \\ \varepsilon & \text{otherwise.} \end{cases}$$

4.1.3 Multi-Set Signatures and Algebras

For convenience, we shall introduce a special kind of signatures and algebras based on multi-sets. First we shall define multi-sets and some basic operations on them.

Definition 4.8 (Multi-Set) A multi-set over a finite non-empty set A is a function

$$\mu : A \rightarrow \mathbb{N}$$

from the set A to the set of natural numbers. For an element $a \in A$, $\mu(a)$ is called the multiplicity of a in μ .

Definition 4.9 (Set of Multi-Sets) The set of all multi-sets over A is denoted by

$$\mathcal{M}(A) = \{\mu \mid \mu : A \rightarrow \mathbb{N}\}.$$

Definition 4.10 (Multi-Set Relations and Operations) Let there be a finite non-empty set A and $\mu_1, \mu_2 \in \mathcal{M}(A)$. We define the following relations and operations:

1. $\mu_1 = \mu_2$ if $\mu_1(a) = \mu_2(a)$ for all $a \in A$ (equality)
2. $\mu_1 \leq \mu_2$ if $\mu_1(a) \leq \mu_2(a)$ for all $a \in A$ (containment)
3. $a \in \mu_1$ if $\mu_1(a) > 0$ (membership)

4. $\mu_1 + \mu_2 = \{\langle a, \mu_1(a) + \mu_2(a) \rangle \mid a \in A\}$ (union)
5. $\mu_1 - \mu_2 = \{\langle a, \max\{0, \mu_1(a) - \mu_2(a)\} \rangle \mid a \in A\}$ (difference)
6. $n \cdot \mu_1 = \{\langle a, n \cdot \mu_1(a) \rangle \mid a \in A\}$, $n \in \mathbb{N}$ (scalar multiplication)
7. $|\mu_1| = \sum_{a \in A} \mu_1(a)$ (cardinality)

A multi-set signature distinguishes two different kinds of sorts: basic sorts and multi-set sorts over basic sorts.

Definition 4.11 (Multi-Set Signature) Let $\mathbf{S} = \langle S, \mathcal{F}, \mathcal{G} \rangle$ be a signature with a finite set of sorts S . Let $S_\beta, S_\mu \subseteq S$ such that $S_\beta \cup S_\mu = S$ and $S_\beta \cap S_\mu = \emptyset$, and let $\mu : S_\beta \rightarrow S_\mu$ be a bijective mapping from basic sorts S_β to multi-set sorts S_μ . Then

$$\mathbf{S}_\mu = \langle S, \mathcal{F}, \mathcal{G}, \mu \rangle$$

is a multi-set signature.

A multi-set algebra is a straightforward extension of an algebra. Based on a multi-set signature, it requires that the carrier of its each multi-set sort is the set of multi-sets over the carrier of the corresponding basic sort.

Definition 4.12 (Multi-Set Algebras) Let \mathbf{S} be a signature and \mathbf{S}_μ be the corresponding multi-set signature. An \mathbf{S} -algebra $\mathcal{A} = \langle \mathcal{D}^\mathcal{A}, \mathcal{F}^\mathcal{A} \rangle$ is an \mathbf{S}_μ -algebra if for each $s \in S_\beta$, $\mathcal{D}_{\mu(s)}^\mathcal{A} = \mathcal{M}(\mathcal{D}_s^\mathcal{A})$.

4.2 Algebraic System Nets

The algebras defined in the previous section form the core of a modelling formalism called Algebraic System Nets [36, 37, 51]. The formalism is a generalisation of Petri Nets [49, 50], which are a kind of generalised automata. Finite state automata have *states* and *actions* leading from one state to another, which are usually represented with circles and labelled directed arcs connecting circles representing states. Only one state is *active* at a time in a finite state automaton.

Petri Nets have *places*, graphically represented with circles, and *transitions*. Unlike the actions in finite state automata, transitions in Petri Nets may connect an arbitrary number of places. A transition, graphically represented with a rectangle, may have a number of *input and output places* connected to it via directed *arcs*. Input places are connected via input arcs (arcs leading to the transition) and output places via output arcs (arcs leading from the transition). Figure 4.1 illustrates the graphical notation typically used with Petri Nets.

In a finite state automaton, one state may be marked active at a time. In a Petri Net, a number of places may be marked with a *token*.² A finite state automaton may take an action if the source state of the action is active. A transition in a Petri Net is *enabled* if all its input places contain enough tokens. Only an enabled

²In Algebraic System Nets, each place may be marked with any number of tokens.

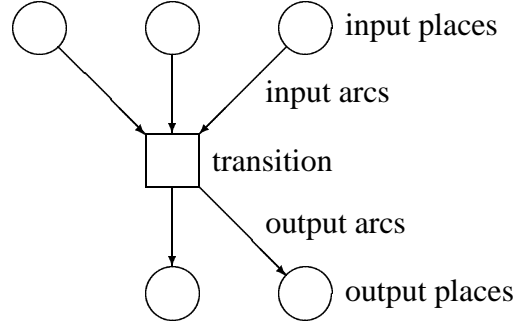


Figure 4.1: Graphical Representation of a Petri Net

transition may *fire*, removing some tokens from its input places and producing some to its output places.³

Petri Nets and Algebraic System Nets clearly have more expressive power than finite state automata. Once we have defined the formal semantics of Algebraic System Nets, we shall introduce the concept of a *reachability graph*, an automaton or a labelled transition system representing the complete behaviour of an Algebraic System Net.

Definition 4.13 (Algebraic System Nets) An algebraic system net

$$\Sigma = \langle \mathcal{N}, \mathcal{A}, \mathcal{V}, i \rangle$$

over \mathcal{A} consists of

1. a net $\mathcal{N} = \langle \mathcal{P}, \mathcal{T}, \mathcal{F} \rangle$ where
 - (a) \mathcal{P} , a finite pairwise disjoint family of sort-indexed places $\mathcal{P} = \bigcup_{s \in S_\mu} \mathcal{P}_s$, is a set of \mathbf{S}_μ -variables whose assignments are multi-set-valued
 - (b) \mathcal{T} , a finite set of transitions, is disjoint from the family of places: $\mathcal{T} \cap \mathcal{P} = \emptyset$
 - (c) $\mathcal{F} \subseteq (\mathcal{T} \times \mathcal{P}) \cup (\mathcal{P} \times \mathcal{T})$ is a flow relation; the items of \mathcal{F} are called arcs; for $t \in \mathcal{T}$ and $p \in \mathcal{P}$, $\langle p, t \rangle \in \mathcal{F}$ is an input arc and $\langle t, p \rangle \in \mathcal{F}$ is an output arc
2. an \mathbf{S}_μ -algebra \mathcal{A} for a multi-set signature $\mathbf{S}_\mu = \langle S_\beta \cup S_\mu, \mathcal{F}, \mathcal{G}, \mu \rangle$; one basic sort $b \in S_\beta$ is the Boolean sort of truth values with $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$
3. a sorted \mathbf{S}_μ -variable set $\mathcal{V} = \bigcup_{s \in S} \mathcal{V}_s$ such that $\mathcal{P}_s \cap \mathcal{V}_s = \emptyset$ for all $s \in S_\mu$
4. a net inscription $i : (\mathcal{P} \cup \mathcal{T} \cup \mathcal{F}) \rightarrow \bigcup_{s \in S} \mathbf{T}_s^{\mathbf{S}_\mu}(\mathcal{V})$ such that
 - (a) $i(p) \in \mathbf{T}_s^{\mathbf{S}_\mu}(\emptyset)$ for each $s \in S_\mu$ and $p \in \mathcal{P}_s$; i.e. $\bigcup_{p \in \mathcal{P}} i(p)$ are the initialisation expressions

³In Algebraic System Nets, tokens are elements of a multi-set; in low-level Petri Nets, tokens are not associated with values.

- (b) $i(t) \in \mathbf{T}_b^{\mathcal{S}_\mu}(\mathcal{V})$ for each $t \in \mathcal{T}$, where the b is the Boolean sort; i.e. $\bigcup_{t \in \mathcal{T}} i(t)$ are the transition guards
- (c) $i(f) \in \mathbf{T}_s^{\mathcal{S}_\mu}(\mathcal{V})$ for each $f \in \mathcal{F}$ such that $f = \langle p, t \rangle$ or $f = \langle t, p \rangle$ where $s \in \mathcal{S}_\mu$, $p \in \mathcal{P}_s$ and $t \in \mathcal{T}$; i.e. $\bigcup_{f \in \mathcal{F}} i(f)$ are the arc expressions.

So far, we have formally defined a structure for modelling concurrent systems. We shall now give the structure a dynamic semantics, which can be applied to determine all possible executions (state–transition sequences) of the model. First we will define the states of the model.

Definition 4.14 (Marking, Local Marking and Tokens) A marking of an algebraic system net Σ is a mapping

$$M : \bigcup_{s \in \mathcal{S}_\mu} (\mathcal{P}_s \rightarrow \mathcal{D}_s^{\mathcal{A}}),$$

and the set of all such markings M is denoted by \mathcal{M}_Σ .

For each $s \in \mathcal{S}_\beta$ and $p \in \mathcal{P}_{\mu(s)}$, we call $M(p)$ the local marking of place p and the items $d \in \mathcal{D}_s^{\mathcal{A}}$ tokens of sort s . We say that the place p contains n tokens carrying the value d if $M(p)(d) = n$.

Definition 4.15 (Initial Marking) Let $v_0 \in V^{\mathcal{A}}(\emptyset)$ be an empty assignment. The marking $M_0 : \mathcal{P} \rightarrow \mathcal{D}^{\mathcal{A}}$ with

$$M_0 : \bigcup_{s \in \mathcal{S}_\mu} \left\{ \left\langle p, e_{v_0}^{\mathcal{A}}(i(p)) \right\rangle \parallel p \in \mathcal{P}_s \right\}$$

is called the initial marking of Σ . Note that the initial marking is undefined if $e_{v_0}^{\mathcal{A}}(i(p)) = \varepsilon$ for some $p \in \mathcal{P}_s$.

The rôle of the places of an algebraic system net has now been covered. The places are associated with the global state of the model, the marking. In each state M of a net Σ , a place $p \in \mathcal{P}$ of Σ contains a multi-set $M(p)$ of tokens. The initial marking yields the initial distribution of tokens in the model, its initial state.

In an algebraic system net, places are connected via arcs and transitions. They define the behaviour of the system, the possible actions leading from one global state to another. In the following definitions, the possible actions in a state are referred to as the enabled transition modes in a marking. Furthermore, a transition can be fired in a marking in an enabled mode to transform the marking to another marking representing another state of the system.

Definition 4.16 (Input and Output Effect) Let Σ be an algebraic system net, $t \in \mathcal{T}$ a transition and $\hat{v}, v \in V^{\mathcal{A}}(\mathcal{V})$ assignments such that v augments \hat{v} , i.e. for each $s \in \mathcal{S}$ and $x \in \mathcal{V}_s$, $\hat{v}(x) = \varepsilon$ or $\hat{v}(x) = v(x)$. The two substitutions $t_{\hat{v}}^-, t_v^+ : \bigcup_{s \in \mathcal{S}_\beta} (\mathcal{P}_{\mu(s)} \rightarrow \check{\mathcal{D}}_{\mu(s)}^{\mathcal{A}})$ are called the input effect and the output effect, respectively, and they are defined by:

$$t_{\hat{v}}^-(p) = \begin{cases} e_{\hat{v}}^{\mathcal{A}}(i(\langle p, t \rangle)) & \text{if } \langle p, t \rangle \in \mathcal{F} \\ \mathcal{D}_s^{\mathcal{A}} \rightarrow \{0\} & \text{otherwise} \end{cases} \quad t_v^+(p) = \begin{cases} e_v^{\mathcal{A}}(i(\langle t, p \rangle)) & \text{if } \langle t, p \rangle \in \mathcal{F} \\ \mathcal{D}_s^{\mathcal{A}} \rightarrow \{0\} & \text{otherwise} \end{cases}$$

Definition 4.17 (Pre-Enabling Rule) Let Σ be an algebraic system net, M its marking, $t \in \mathcal{T}$ a transition of Σ and $\hat{v} \in V^{\mathcal{A}}(\mathcal{V})$ an assignment. Transition t is pre-enabled in mode \hat{v} at marking M of Σ if the following conditions hold for each $s \in \mathcal{S}_\mu$ and $p \in \mathcal{P}_s$:

1. $e_{\hat{v}}^{\mathcal{A}}(i(t)) = \top$; i.e. the transition guard holds
2. $t_{\hat{v}}^-(p) \neq \varepsilon$; i.e. the input effect is defined
3. $t_{\hat{v}}^-(p) \leq M(p)$; i.e. each place contains enough tokens

Definition 4.18 (Enabling Rule) Let Σ , M , $t \in \mathcal{T}$ and $\hat{v} \in V^{\mathcal{A}}(\mathcal{V})$ be such that t is pre-enabled in mode \hat{v} at marking M of Σ . Let $v \in V^{\mathcal{A}}(\mathcal{V})$ such that for each $s \in \mathcal{S}$ and $x \in \mathcal{V}_s$, $\hat{v}(x) = \varepsilon$ or $\hat{v}(x) = v(x)$. Transition t is enabled in mode v if for each $s \in \mathcal{S}_\mu$ and $p \in \mathcal{P}_s$ it holds that $t_v^+(p) \neq \varepsilon$.

Our definition of the transition enabling rule is divided into two parts, one based on the input effect and another one based on the output effect. Our definition distinguishes a set of variables

$$\mathcal{V}_o = \bigcup_{s \in \mathcal{S}} \{x \in \mathcal{V}_s \mid \hat{v}(x) \neq v(x)\}$$

that are defined (not ε) in v but undefined in \hat{v} . These variables are called *output variables*, since they can only be evaluated on the output arcs of an enabled transition. An assignment⁴ \hat{v} can be extended to a number of assignments v e.g. by enumerating through the domain of each variable in \mathcal{V}_o , picking values such that a user-defined condition $c \in \mathbf{T}_b^{\mathcal{S}_\mu}(\mathcal{V})$ holds: $e_v^{\mathcal{A}}(c) = \top$. The implementation in [41] does this, and it also provides variables for $t_{\hat{v}}^-(p)$ for each place p .⁵ Formally, for each sort $s \in \mathcal{S}_\mu$ and place $p \in \mathcal{P}_s$, we have a variable $x_p \in \mathcal{V}_s$, and for each enabled instance $\hat{v}, v \in V^{\mathcal{A}}(\mathcal{V})$ of a transition $t \in \mathcal{T}$, it holds that

$$\begin{aligned} \hat{v}(x_p) &= \varepsilon \\ v(x_p) &= t_{\hat{v}}^-(p). \end{aligned}$$

Note that Definition 4.13 does not allow variables to refer to the global state of the model, which would make it possible to simulate Turing machines [46, Chapter 2] with Algebraic System Nets.⁶

Definition 4.19 (Firing Rule) Let Σ , M , $t \in \mathcal{T}$ and $\hat{v}, v \in V^{\mathcal{A}}(\mathcal{V})$ be such that t is enabled in mode v at marking M of Σ . The firing of transition t in mode v at marking M produces a marking

$$M' = \bigcup_{s \in \mathcal{S}_\mu} \{ \langle p, M(p) - t_{\hat{v}}^-(p) + t_v^+(p) \rangle \mid p \in \mathcal{P}_s \}.$$

⁴Also the words “instance” and “valuation” are usual.

⁵This extension to the formalism is redundant in the sense that a model making use of output variables can be transformed to a model without output variables. However, using output variables will speed up the reachability analysis and can make models more intuitive.

⁶As one of the consequences, the reachability problem for Algebraic System Nets would become undecidable, since an Algebraic System Net could decide whether a Turing machine halts.

The fact that M' is the result of firing t in mode v at M can be written $M[t_v]M'$.

The firing rule allows us to determine the set of markings that are reachable in the model, the reachable state space of the model.

Definition 4.20 (Reachable States) Let Σ be an algebraic system net and M_0 the initial marking of Σ . The set of reachable states of Σ is the smallest set $R \subseteq \mathcal{M}_\Sigma$ fulfilling the following conditions:

1. $M_0 \in R$
2. $\{M' \parallel M[t_v]M'\} \subseteq R$ for all $M \in R$, $t \in \mathcal{T}$ and $v \in V^{\mathcal{A}}(\mathcal{V})$ such that t is enabled in mode v at marking M .

Definition 4.21 (Reachable Actions) Let Σ be an algebraic system net and M_0 the initial marking of Σ and R the set of reachable states of Σ . The set of reachable actions of Σ is defined to be the smallest set $E \subseteq R \times (\mathcal{T} \times V^{\mathcal{A}}(\mathcal{V})) \times R$ for which

$$\{\langle M, \langle t, v \rangle, M' \rangle \parallel M[t_v]M'\} \subseteq E$$

for all $M \in R$, $t \in \mathcal{T}$ and $v \in V^{\mathcal{A}}(\mathcal{V})$ such that t is enabled in mode v at M .

Definition 4.22 (Reachability Graph) Let Σ be an algebraic system net and M_0 the initial marking of Σ . Let R be the set of reachable states and E the set of reachable actions of Σ . The reachability graph of Σ is the directed graph

$$G = \langle R, E \rangle.$$

Chapter 5

Data Types

Digital devices such as computers represent all data with binary digits, also known as bits, of zeros and ones. They are often grouped to fixed-width *machine words* of $n = 2^m$ bits, often interpreted as integer numbers between 0 and $2^n - 1$ or between -2^{n-1} and $2^{n-1} - 1$. This is fine for numerical applications, but many other applications would benefit from more sophisticated structures for managing data. The data model of practically all high-level programming languages is based on *data types*.

Mathematically, a data type can be represented as a set comprising all the acceptable values, the *domain* of the type. A computer implementation of data types has to be more specific. We will present an inductive definition of a data type system that has been implemented in [41].

5.1 Design Criteria

Our data type system is based on the following design criteria.

Limited domain

All data types \mathcal{D} have a limited domain, $|\mathcal{D}| > 0$, facilitating a conversion between data items and sequences of machine words. The domains can be limited further by specifying ranges of allowed values.

Total order

For all data types \mathcal{D} and for all data items $i, j \in \mathcal{D}$, there is an asymmetric irreflexive transitive relation $<_{\mathcal{D}} \subset \mathcal{D} \times \mathcal{D}$,¹ i.e. at most one of the following holds: $i <_{\mathcal{D}} j$ or $j <_{\mathcal{D}} i$. If neither property holds, i and j refer to the same data item: $i = j$.

Tight representation

For a data type \mathcal{D} with a domain consisting of $n = |\mathcal{D}|$ data items, we construct a bijective mapping $o_{\mathcal{D}} : \mathcal{D} \rightarrow \{0, \dots, n-1\}$ such that for all $d, d' \in \mathcal{D}$, $o_{\mathcal{D}}(d) < o_{\mathcal{D}}(d')$ if and only if $d <_{\mathcal{D}} d'$. The mapping allows each $d \in \mathcal{D}$ to be represented with $\lceil \log_2 n \rceil$ binary digits.

¹We deviate from the commonly used definition that includes reflexivity, i.e. $i \leq_{\mathcal{D}} i$.

Expressive power

Most data types used in programming languages should have a straightforward transformation to our data types. We have to omit pointers to data items, since they would violate all the preceding properties, and unbounded data types such as lists and trees. We do have variable-length buffers of limited capacity.

5.1.1 Tight Representation

For a data type \mathcal{D} with $|\mathcal{D}| = n$ and with a total order $<_{\mathcal{D}} \subset \mathcal{D} \times \mathcal{D}$, we define the mapping

$$o_{\mathcal{D}} : \mathcal{D} \rightarrow \{0, \dots, n-1\} : d \mapsto |\{k \in \mathcal{D} \mid k <_{\mathcal{D}} d\}|.$$

It is easy to see that the mapping is bijective and that it preserves the order of the mapped items. Because $<_{\mathcal{D}}$ is a total order, \mathcal{D} can be written as

$$\mathcal{D} = \{d_0, \dots, d_{n-1}\}$$

such that $d_{i-1} <_{\mathcal{D}} d_i$ for all $0 < i < n$. Now $o_{\mathcal{D}}$ maps each d_i , $0 \leq i < n$, to a unique value:

$$\begin{aligned} o_{\mathcal{D}}(d_i) &= |\{k \in \mathcal{D} \mid k <_{\mathcal{D}} d_i\}| \\ &= |\{d_0, \dots, d_{i-1}\}| \\ &= i. \end{aligned}$$

Since $o_{\mathcal{D}}(d_i) = i$, it holds that $o_{\mathcal{D}}(d_i) < o_{\mathcal{D}}(d_j)$ if and only if $i < j$, or $d_i <_{\mathcal{D}} d_j$. Thus, $o_{\mathcal{D}}$ is an order-preserving mapping.

5.1.2 Expressive Power

Pointers to data items are a problematic issue. We cannot allow them for several reasons. First of all, pointers having a large number of possible addressees (data objects they may point to) would destroy the modularity and locality properties of the verification model, thus restricting the applicability of compositional verification techniques. Second, pointers do not mix with the fundamental concept of Petri Nets: transitions affect the state only locally. Furthermore, pointers can be used to define data types of unlimited domain, which cannot be handled by analysis techniques that require unfolding; see Section 7.1.2. Last but not least, no total order can be defined for pointers in an obvious way independent of the underlying computer system.

Some restricted cases of pointer usage can be modelled by using the identifier type that will be introduced in Section 5.2.5; this will be discussed in Section 9.3.2.

5.2 Simple Types

There are a number of data types whose data items can be directly represented with machine words. Usually machine words are interpreted as integers, but they can easily be interpreted as other enumerable unstructured data items as well.

5.2.1 Boolean

The Boolean type $\mathbb{B} = \{\perp, \top\}$ represents truth values. It has the total order $<_{\mathbb{B}} = \{\langle \perp, \top \rangle\}$.

5.2.2 Character

The Character type \mathbb{K} represents the character set of the underlying computer system. The size of the alphabet is typically $|\mathbb{K}| = 2^8 = 256$. A bijective mapping

$$c : \mathbb{K} \rightarrow \{0, \dots, |\mathbb{K}| - 1\}$$

assigns the character set an ordering. The total order $<_{\mathbb{K}} \subset \mathbb{K} \times \mathbb{K}$ is defined as

$$<_{\mathbb{K}} = \{\langle i, j \rangle \in \mathbb{K} \times \mathbb{K} \mid c(i) < c(j)\}$$

based on the bijective mapping c and the total order in the set of integer numbers.

5.2.3 Integer

The Integer type \mathbb{I} is a straightforward representation of machine words. Machine words can be interpreted either as signed or as unsigned integers.² The domain of the signed variant usually is $\{-2^{n-1}, \dots, 2^{n-1} - 1\}$, where n is the length of the machine word in bits, typically 2^5 or 2^6 . The total order $<_{\mathbb{I}}$ is defined in the obvious way:

$$<_{\mathbb{I}} = \{\langle i, j \rangle \in \mathbb{I} \times \mathbb{I} \mid i < j\}.$$

5.2.4 Enumerated Types

Enumerated data types are similar to integers. An enumerated type consists of named integer constants

$$\mathbb{E}_N = \{\langle k, n_k \rangle \mid k \in N \wedge n_k \in \mathbb{I}\}$$

having distinct values:

$$\forall k \in N : \forall l \in N : n_k = n_l \Rightarrow k = l.$$

The total order $<_{\mathbb{E}_N}$ is defined in terms of the integer values of the named constants:

$$<_{\mathbb{E}_N} = \{\langle \langle i, n_i \rangle, \langle j, n_j \rangle \rangle \in \mathbb{E}_N \times \mathbb{E}_N \mid n_i <_{\mathbb{I}} n_j\}.$$

²Like many programming languages, our implementation [41] provides both unsigned and signed integers, but for the sake of simplicity, we make the assumption that all integers are of the same type.

5.2.5 Identifier Type

The identifier type is comparable to a pointer or a resource handle in a programming language environment. From the user's point of view, the type does not have any literals (constants), and identifier values can only be compared for equality; for the user, there is no total order or conversion to integers. Symmetry reductions are very effective on models making use of this type.

The identifier type \mathbb{J}_n is represented with integers,

$$\mathbb{J}_n = \{0, \dots, n-1\},$$

and the total order $<_{\mathbb{J}_n}$ is defined in the obvious way:

$$<_{\mathbb{J}_n} = \{\langle i, j \rangle \in \mathbb{J}_n \times \mathbb{J}_n \mid i < j\}.$$

5.3 Structured Types

Structured data types add expressive power to the data type system, especially when structured types can be constructed of other structured types without limitation. As all types in our data type system must have finite domains, we cannot allow any kind of recursion in data type definitions. For instance, it is impossible to define a tuple A with a component belonging to a union type containing the tuple A .

5.3.1 Tuple

A tuple data type is a Cartesian product of data types \mathcal{D}_k ,

$$\mathbb{T}_{\mathcal{D}_1 \dots \mathcal{D}_n} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n.$$

Syntactically, the data types \mathcal{D}_k in a tuple are called the *components* of the tuple. Typically components are identified by names; our formal description shall use index numbers:

$$f_{\mathbb{T}_{\mathcal{D}_1 \dots \mathcal{D}_n}} : \mathbb{T}_{\mathcal{D}_1 \dots \mathcal{D}_n} \times \{1, \dots, n\} \rightarrow \bigcup_{k=1}^n \mathcal{D}_k : \langle \langle d_1, \dots, d_n \rangle, k \rangle \mapsto d_k$$

The total order $<_{\mathcal{D}}$ of $\mathcal{D} = \mathbb{T}_{\mathcal{D}_1 \dots \mathcal{D}_n}$ is defined lexicographically in the little endian [5] way: The first component is the least significant one, and it will thus determine the ordering only when all other components are equal:

$$<_{\mathcal{D}} = \left\{ \langle \langle i_1, \dots, i_n \rangle, \langle j_1, \dots, j_n \rangle \rangle \in \mathcal{D} \times \mathcal{D} \left\| \bigvee_{k=1}^n (i_k <_{\mathcal{D}_k} j_k \wedge \bigwedge_{l=k+1}^n i_l = j_l) \right. \right\}.$$

When $n = 0$, $\mathcal{D} = \mathbb{T} = \{\langle \rangle\}$ and $<_{\mathcal{D}} = \emptyset$, as we define the empty disjunction to be false.

Why is $<_{\mathcal{D}}$ a total order? For $l > k$, $i_l = j_l$ implies that $i_l <_{\mathcal{D}_l} j_l$ cannot hold (since $<_{\mathcal{D}_l}$ is a total order). On the other hand, if $i_k <_{\mathcal{D}_k} j_k$, we cannot have $i_k = j_k$. Therefore, at most one of the disjuncts in the condition can be true at a time. We shall now show that for any $i, j \in \mathcal{D}$, either $i <_{\mathcal{D}} j$, $j <_{\mathcal{D}} i$ or $i = j$. There are two cases. Either one disjunct k' is true or all disjuncts are false. When a disjunct k' is true, we have that for $l > k'$, $i_l = j_l$, and that $i_{k'} <_{\mathcal{D}_{k'}} j_{k'}$. The symmetric comparison yields $j_l = i_l$ for $l > k'$, but neither $j_{k'} <_{\mathcal{D}_{k'}} i_{k'}$ nor $j_{k'} = i_{k'}$ holds, because $<_{\mathcal{D}_{k'}}$ is a total order. Therefore $j <_{\mathcal{D}} i$ does not hold in this case. When all disjuncts in $i <_{\mathcal{D}} j$ are false, either $i_k = j_k$ for all $1 \leq k \leq n$, in which case $i = j$, or there exists a k'' such that $j_{k''} <_{\mathcal{D}_{k''}} i_{k''}$ and thus $j <_{\mathcal{D}} i$.

5.3.2 Associative Array

An associative array is a collection of elements of a data type \mathcal{D}_e indexed by data items of \mathcal{D}_x . Formally,

$$\mathbb{A}_{\mathcal{D}_x, \mathcal{D}_e} = \underbrace{\mathcal{D}_e \times \cdots \times \mathcal{D}_e}_{|\mathcal{D}_x| \text{ terms}}$$

Arrays are similar to tuples, except that all components (elements) are of the same type \mathcal{D}_e and that the elements are indexed by integer representations $o_{\mathcal{D}_x}(d)$ of the index values $d \in \mathcal{D}_x$ rather than by component names or numbers:

$$x_{\mathbb{A}_{\mathcal{D}_x, \mathcal{D}_e}} : \mathbb{A}_{\mathcal{D}_x, \mathcal{D}_e} \times \mathcal{D}_x \rightarrow \mathcal{D}_e : \langle \langle d_1, \dots, d_{|\mathcal{D}_x|} \rangle, d \rangle \mapsto d_{1+o_{\mathcal{D}_x}(d)}.$$

The total order $<_{\mathcal{D}}$ of $\mathcal{D} = \mathbb{A}_{\mathcal{D}_x, \mathcal{D}_e}$ is defined lexicographically, in a way similar to the total order among tuples:

$$<_{\mathcal{D}} = \left\{ \left\langle \langle i_1, \dots, i_{|\mathcal{D}_x|} \rangle, \langle j_1, \dots, j_{|\mathcal{D}_x|} \rangle \right\rangle \in \mathcal{D} \times \mathcal{D} \left\| \bigvee_{k=1}^{|\mathcal{D}_x|} (i_k <_{\mathcal{D}_e} j_k \wedge \bigwedge_{l=k+1}^{|\mathcal{D}_x|} i_l = j_l) \right. \right\}.$$

Since the construction is essentially the same as for tuples (substituting $|\mathcal{D}_x|$ for n), it is easy to see that $<_{\mathcal{D}}$ is a total order.

5.3.3 Variable-Length Buffer

Linked lists, first-in-first-out queues and last-in-first-out stacks can be viewed as variable-length buffers of data elements of type \mathcal{D}_e . Since all our data types have a bounded domain, our variable-length buffer has a maximum number n of data elements. Formally,

$$\mathbb{V}_{\mathcal{D}_e, n} = \bigcup_{k=0}^n \mathcal{D}_e^k$$

where the superscript indicates a Cartesian product

$$\mathcal{D}_e^k = \underbrace{\mathcal{D}_e \times \cdots \times \mathcal{D}_e}_{k \text{ times}}$$

The total order $<_{\mathcal{D}}$ of $\mathcal{D} = \mathbb{V}_{\mathcal{D}_e, n}$ is defined by considering a missing item to be the smallest. A buffer value can only contain missing items as a contiguous sequence in its tail. Because of this, the lexicographical comparison can be implemented by comparing the actual buffer length first:

$$<_{\mathcal{D}} = \left\{ \langle \langle (i)_1^p \rangle, \langle (j)_1^q \rangle \rangle \in \mathcal{D} \times \mathcal{D} \mid p < q \vee \left(p = q \wedge \bigvee_{k=1}^m (i_k <_{\mathcal{D}_e} j_k) \wedge \bigwedge_{l=k+1}^m i_l = j_l \right) \right\}$$

where m denotes the minimum of p and q , and the abbreviation $(i)_1^p$ stands for the sequence i_1, \dots, i_p .

5.3.4 Tagged Union

The union type $\mathbb{U}_{\mathcal{D}_1 \dots \mathcal{D}_n}$ can hold values of different types $\mathcal{D}_1, \dots, \mathcal{D}_n$. In order to ease type conversions, a value of type \mathcal{D}_k will be tagged with the index number k .

$$\mathbb{U}_{\mathcal{D}_1 \dots \mathcal{D}_n} = \bigcup_{k=1}^n \{ \langle k, d_k \rangle \mid d_k \in \mathcal{D}_k \}$$

The total order $<_{\mathcal{D}}$ for $\mathcal{D} = \mathbb{U}_{\mathcal{D}_1 \dots \mathcal{D}_n}$ first compares index numbers and then values:

$$<_{\mathcal{D}} = \left\{ \langle \langle k_i, i \rangle, \langle k_j, j \rangle \rangle \in \mathcal{D} \times \mathcal{D} \mid k_i < k_j \vee (k_i = k_j \wedge i <_{\mathcal{D}_{k_i}} j) \right\}.$$

5.4 Constraints

Type constraints limit the domain of a data type. A constraint function f maps each data item of a data type \mathcal{D} to a Boolean value that specifies whether the data item is allowed in the constrained type

$$\mathcal{C}(\mathcal{D}, f : \mathcal{D} \rightarrow \mathbb{B}) = \{ d \in \mathcal{D} \mid f(d) = \top \}.$$

The total order $<_{\mathcal{C}(\mathcal{D}, f)}$ is defined in terms of the total order $<_{\mathcal{D}}$ of the underlying data type \mathcal{D} :

$$<_{\mathcal{C}(\mathcal{D}, f)} = \{ \langle i, j \rangle \in \mathcal{C}(\mathcal{D}, f) \times \mathcal{C}(\mathcal{D}, f) \mid i <_{\mathcal{D}} j \}.$$

In [41], constraints are implemented as lists of non-overlapping semi-open ranges interpreted as unions, and constraints can be defined for all data types. The grammar of the input language allows both unions and intersections of ranges to be entered. Figure 5.1 depicts how unions and intersections of ranges can be evaluated to convert constraints to the canonic form. Symmetric cases (obtained by swapping the rôles of r_1 and r_2) are omitted from the figure.

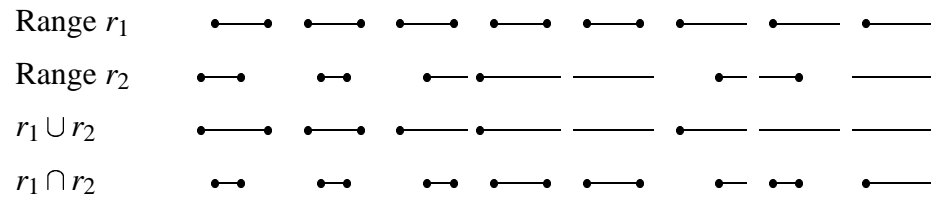


Figure 5.1: The Union and the Intersection of Ranges

5.4.1 Computing the Union

When computing the union of two constraints c_1 and c_2 , the algorithm implemented in [41] compares each range r_i^1 in c_1 with each range r_j^2 in c_2 . If one of the cases depicted in Figure 5.1 applies, then the union of the two ranges will be added to the result. Otherwise the algorithm will compare whether the upper limit of r_i^1 is one less than the lower limit of r_j^2 or vice versa, and add the combination of the two ranges to the result if this is the case. This is how $\{1, 2\} \cup \{3\}$ will become $\{1, 2, 3\}$. If even this does not apply, the two ranges are disjoint, and both will be stored in the resulting constraint.

5.4.2 Computing the Intersection

The intersection of two constraints c_1 and c_2 is computed by comparing each range r_i^1 in c_1 with each range r_j^2 in c_2 . If one of the cases described in Figure 5.1 applies, the intersection of the two ranges will be added to the resulting constraint. Otherwise the two ranges are disjoint, and the resulting constraint will not be augmented.

Chapter 6

Algebraic Operations

Computers work by performing operations on data. Big operations are defined in terms of smaller operations, and at the lowest level, there are *basic operations* performed by the underlying computing machinery. Because operations are performed on data items, they are tightly coupled with the data type system. Here we shall present the basic operations implemented in [41]; the data type system is defined in Chapter 5.

It is quite common that there is some redundancy among basic operations, even in low-level languages. For instance, if the basic operations of an algebra include logical conjunction, logical disjunction and logical negation, the conjunction or the disjunction can be constructed in terms of the two remaining operations by applying De Morgan's law: $a \wedge b = \neg(a \vee b)$. High-level languages tend to have more redundant operations than low-level languages, so called *syntactic sugar* that adds expressive power to the language and allows for compact notation.

High-level programming languages define expressions and statements. In the Algebraic System Nets defined in Chapter 4, expressions do not have side effects (they cannot affect the assignment they are evaluated in), and there is only one form of a statement, the transition. A transition can be viewed as a statement that changes the value of some variables in an assignment.¹

The algebraic operations defined in this chapter use the notational conventions

$$\begin{aligned} \mathbf{op} &\in \mathcal{F}_{s_1, \dots, s_n, s} \\ \mathbf{op}^{\mathcal{A}} &: \check{\mathcal{D}}_{s_1}^{\mathcal{A}} \times \dots \times \check{\mathcal{D}}_{s_n}^{\mathcal{A}} \rightarrow \check{\mathcal{D}}_s^{\mathcal{A}} \end{aligned}$$

and the implicit definition

$$\mathbf{op}^{\mathcal{A}}(d) = \varepsilon \quad \text{for all } d \in \left(\check{\mathcal{D}}_{s_1}^{\mathcal{A}} \times \dots \times \check{\mathcal{D}}_{s_n}^{\mathcal{A}} \right) \setminus \left(\mathcal{D}_{s_1}^{\mathcal{A}} \times \dots \times \mathcal{D}_{s_n}^{\mathcal{A}} \right).$$

Further algebraic operation definitions in this chapter augment the implicit definition. In further definitions, the domains consist of the carriers of the domain sorts (excluding the symbol ε), and the ranges are the carriers of the range sorts augmented with the symbol ε .

¹In Algebraic System Nets, these variables (in the algebraic sense) are the local markings of the places attached to the transition via input and output arcs.

6.1 Design Criteria

The basic algebraic operations were chosen according to the following criteria:

No side effects

No operation modifies the environment it is evaluated in.

Total order

With the exception of the identifier type defined in Section 5.2.5, there are operations for accessing the total order of data items e.g. via comparisons and via predecessor and successor operations.

Expressive power

All basic operations of popular programming languages such as C [28] are covered, except for pointer arithmetics and operations having side effects. In addition, there are some operations that deal with multi-sets, which are an extension over the data type system defined in Chapter 5.

6.2 Variables

In computer languages, there typically are two things that can be done to variables: they can be *declared* and *referenced*. In various high-level Petri Net formalisms, variables are often declared implicitly, at the points when the variables are first referenced. Declaring variables corresponds to extending the family of variables $\mathcal{V} = \bigcup_{s \in \mathcal{S}} \mathcal{V}_s$ with new items.

Variable references $x \in \mathcal{V}_s$ for $s \in \mathcal{S}$ are not operations, but direct members of the set of terms: $x \in \mathbf{T}_s^{\mathcal{S}}(\mathcal{V})$.

6.3 Operations on Basic Sorts

The lowest level operations of our algebra are defined for the family of basic sorts \mathcal{S}_β , which is mapped to data types constructed using the structures defined in Chapter 5.

6.3.1 Constants

Constants in our algebra include literals of simple types, minimum and maximum value of ordered types, and integer constants representing the number of elements in a type. For each basic sort $s \in \mathcal{S}_\beta$ and $d \in \mathcal{D}_s^{\mathcal{A}}$, we define

$$\begin{aligned} \mathbf{constant}_d &\in \mathcal{F}_{\lambda, s} \\ \mathbf{constant}_d^{\mathcal{A}}() &= d. \end{aligned}$$

Note that we do not fix the syntax here. Especially if an implementation performs *constant folding*, replacing terms $g \in \mathbf{T}_s^{\mathcal{S}}(\emptyset)$ with terms $\mathbf{constant}_d^{\mathcal{A}}()$ where

$d = e_v^{\mathcal{A}}(g)$ for an empty assignment $v \in V^{\mathcal{A}}(\emptyset)$, there are numerous ways of writing constants.

There is also an undefined constant, whose purpose is to detect errors in an Algebraic System Net model. Remember, a transition instance can only be fired when all its arc expressions evaluate to defined values. In [41], there are two variations of the undefined constant: one that only causes an error message to be displayed for the current state, and another that causes the reachability analysis to be aborted.

$$\begin{aligned} \mathbf{undefined} &\in \mathcal{F}_{\lambda, s} \\ \mathbf{undefined}^{\mathcal{A}}() &= \varepsilon. \end{aligned}$$

6.3.2 Total Order

Successor and Predecessor

For each sort $s \in \mathcal{S}_{\beta}$ with the carrier $\mathcal{D}_s^{\mathcal{A}}$ and the bijective mapping $o_{\mathcal{D}_s^{\mathcal{A}}} : \mathcal{D}_s^{\mathcal{A}} \rightarrow \{0, \dots, |\mathcal{D}_s^{\mathcal{A}}| - 1\}$ defined in Section 5.1.1, we define the *successor* and *predecessor* operations as follows:

$$\begin{aligned} \mathbf{succ} &\in \mathcal{F}_{s, s} \\ \mathbf{succ}^{\mathcal{A}}(d) &= d_+ \\ \mathbf{pred} &\in \mathcal{F}_{s, s} \\ \mathbf{pred}^{\mathcal{A}}(d) &= d_- \end{aligned}$$

such that

$$\begin{aligned} o_{\mathcal{D}_s^{\mathcal{A}}}(d) + 1 &\equiv o_{\mathcal{D}_s^{\mathcal{A}}}(d_+) \pmod{|\mathcal{D}_s^{\mathcal{A}}|} \\ o_{\mathcal{D}_s^{\mathcal{A}}}(d) - 1 &\equiv o_{\mathcal{D}_s^{\mathcal{A}}}(d_-) \pmod{|\mathcal{D}_s^{\mathcal{A}}|}. \end{aligned}$$

Comparison

For the Boolean sort $b \in \mathcal{S}_{\beta}$ with the carrier $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$ and for each sort $s \in \mathcal{S}_{\beta}$ with the carrier $\mathcal{D}_s^{\mathcal{A}}$ and the total order $<_{\mathcal{D}_s^{\mathcal{A}}} \subset \mathcal{D}_s^{\mathcal{A}} \times \mathcal{D}_s^{\mathcal{A}}$, we define following comparison operations:

$$\begin{aligned} \mathbf{equal} &\in \mathcal{F}_{s, s, b} \\ \mathbf{unequal} &\in \mathcal{F}_{s, s, b} \\ \mathbf{less} &\in \mathcal{F}_{s, s, b} \\ \mathbf{greater} &\in \mathcal{F}_{s, s, b} \\ \mathbf{lessequal} &\in \mathcal{F}_{s, s, b} \\ \mathbf{greaterequal} &\in \mathcal{F}_{s, s, b} \end{aligned}$$

$$\begin{aligned}
\mathbf{equal}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_1 = d_2 \\ \perp & \text{otherwise} \end{cases} \\
\mathbf{unequal}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \perp & \text{if } d_1 = d_2 \\ \top & \text{otherwise} \end{cases} \\
\mathbf{less}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_1 <_{\mathcal{D}_s^{\mathcal{A}}} d_2 \\ \perp & \text{otherwise} \end{cases} \\
\mathbf{greater}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_2 <_{\mathcal{D}_s^{\mathcal{A}}} d_1 \\ \perp & \text{otherwise} \end{cases} \\
\mathbf{lessequal}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \perp & \text{if } d_2 <_{\mathcal{D}_s^{\mathcal{A}}} d_1 \\ \top & \text{otherwise} \end{cases} \\
\mathbf{greaterequal}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \perp & \text{if } d_1 <_{\mathcal{D}_s^{\mathcal{A}}} d_2 \\ \top & \text{otherwise} \end{cases}
\end{aligned}$$

Note that equality and inequality comparisons are defined for all basic sorts. The implementation in [41] restricts the use of the other comparison operations on sorts that do not contain the identifier sort.

6.3.3 Logical Operations

For the Boolean sort $b \in \mathcal{S}_\beta$ with the carrier $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$, we define the following logical operations:

$$\begin{aligned}
\mathbf{not} &\in \mathcal{F}_{b,b} \\
\mathbf{and} &\in \mathcal{F}_{b,b,b} \\
\mathbf{or} &\in \mathcal{F}_{b,b,b} \\
\mathbf{impl} &\in \mathcal{F}_{b,b,b} \\
\mathbf{xor} &\in \mathcal{F}_{b,b,b} \\
\mathbf{equiv} &\in \mathcal{F}_{b,b,b} \\
\mathbf{not}^{\mathcal{A}}(d_1) &= \begin{cases} \perp & \text{if } d_1 = \top \\ \top & \text{otherwise} \end{cases} \\
\mathbf{and}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_1 = \top \text{ and } d_2 = \top \\ \perp & \text{otherwise} \end{cases} \\
\mathbf{or}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_1 = \top \text{ or } d_2 = \top \\ \perp & \text{otherwise} \end{cases} \\
\mathbf{impl}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_1 = \perp \text{ or } d_2 = \top \\ \perp & \text{otherwise} \end{cases} \\
\mathbf{xor}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \perp & \text{if } d_1 = d_2 \\ \top & \text{otherwise} \end{cases} \\
\mathbf{equiv}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \top & \text{if } d_1 = d_2 \\ \perp & \text{otherwise} \end{cases}
\end{aligned}$$

6.3.4 Integer Arithmetics

For the integer sort $s \in \mathcal{S}_\beta$ with the carrier $\mathcal{D}_s^{\mathcal{A}} = \mathbb{I}$, we define the following basic arithmetic operations based on addition, subtraction, multiplication and truncated division of integer numbers. To simplify the following definitions, we do not explicitly state that when the result of an integer arithmetic operation does not belong to \mathbb{I} , it will be ε . Also, our definition of the division operation rounds the values towards negative infinity. Our implementation [41] uses the underlying C++ [27] implementation, which can as well round the values towards zero. To reduce implementation-defined behaviour, our modulus operation is only defined for positive integers.

$$\begin{aligned}
 \mathbf{negate} &\in \mathcal{F}_{s,s} \\
 \mathbf{plus} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{minus} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{times} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{divide} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{modulus} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{negate}^{\mathcal{A}}(d_1) &= 0 - d_1 \\
 \mathbf{plus}^{\mathcal{A}}(d_1, d_2) &= d_1 + d_2 \\
 \mathbf{minus}^{\mathcal{A}}(d_1, d_2) &= d_1 - d_2 \\
 \mathbf{times}^{\mathcal{A}}(d_1, d_2) &= d_1 \cdot d_2 \\
 \mathbf{divide}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \varepsilon & \text{if } d_2 = 0 \\ \lfloor \frac{d_1}{d_2} \rfloor & \text{otherwise} \end{cases} \\
 \mathbf{modulus}^{\mathcal{A}}(d_1, d_2) &= \begin{cases} \varepsilon & \text{if } d_1 < 0 \\ \varepsilon & \text{if } d_2 \leq 0 \\ d_1 - d_2 \cdot \lfloor \frac{d_1}{d_2} \rfloor & \text{otherwise} \end{cases}
 \end{aligned}$$

Furthermore, in [41] we define the following binary arithmetic operations for integers. Their definition depends on the underlying implementation of the built-in C++ [27] operators \sim , $\&$, $|$, \wedge , \ll and \gg , respectively. The bitwise logical operators will be evaluated without errors; the shift operators will evaluate to ε if the second argument, denoting the amount of bits to be shifted, is negative or at least the width of the machine word.

$$\begin{aligned}
 \mathbf{bitnot} &\in \mathcal{F}_{s,s} \\
 \mathbf{bitand} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{bitor} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{bitxor} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{shiffl} &\in \mathcal{F}_{s,s,s} \\
 \mathbf{shiftr} &\in \mathcal{F}_{s,s,s}
 \end{aligned}$$

6.3.5 Structure Operations

Structured data types, defined in Section 5.3, can be viewed as containers for data items. The operations defined for structured data types mostly deal with composing and decomposing data items of structured types.

Tuple

The tuple type, defined in Section 5.3.1, has two operations: composition and decomposition. Formally, for all $s, s_1, \dots, s_n \in \mathcal{S}_\beta$ such that $\mathcal{D}_s^{\mathcal{A}} = \mathbb{T}_{\mathcal{D}_{s_1}^{\mathcal{A}} \dots \mathcal{D}_{s_n}^{\mathcal{A}}}$ and for $1 \leq k \leq n$, we define the following operations:

$$\begin{aligned} \mathbf{cons}_s &\in \mathcal{F}_{s_1, \dots, s_n, s} \\ \mathbf{component}_{s,k} &\in \mathcal{F}_{s, s_k} \\ \mathbf{cons}_s^{\mathcal{A}}(d_1, \dots, d_n) &= \langle d_1, \dots, d_n \rangle \\ \mathbf{component}_{s,k}^{\mathcal{A}}(\langle d_1, \dots, d_n \rangle) &= d_k \end{aligned}$$

Associative Array

For the two variants of the associative array defined in Section 5.3.2, there are two basic operations, which very closely resemble the basic operations defined for tuples.

For all $s, s_1, s_2 \in \mathcal{S}_\beta$ such that $\mathcal{D}_s^{\mathcal{A}} = \mathbb{A}_{\mathcal{D}_{s_1}^{\mathcal{A}}, \mathcal{D}_{s_2}^{\mathcal{A}}}$ and for $n = |\mathcal{D}_{s_1}^{\mathcal{A}}|$, we define the following basic operations:

$$\begin{aligned} \mathbf{cons}_s &\in \mathcal{F}_{s_2^n, s} \\ \mathbf{index}_s &\in \mathcal{F}_{s, s_1, s_2} \\ \mathbf{cons}_s^{\mathcal{A}}(d_1, \dots, d_n) &= \langle d_1, \dots, d_n \rangle \\ \mathbf{index}_s^{\mathcal{A}}(\langle d_1, \dots, d_n \rangle, d) &= d_k \quad \text{such that } k = 1 + o_{\mathcal{D}_{s_1}^{\mathcal{A}}}(d) \end{aligned}$$

where s_2^n stands for $\underbrace{s_2, \dots, s_2}_{n \text{ times}}$.

Variable-Length Buffer

For the variable-length buffer defined in Section 5.3.3, our basic operations extend the set of operations traditionally defined for queues and stacks. For all $s, s_1, s' \in \mathcal{S}_\beta$ and $n \geq 0$ such that $\mathcal{D}_s^{\mathcal{A}} = \mathbb{V}_{\mathcal{D}_{s_1}^{\mathcal{A}}, n}$ and $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{I}$, and for $0 \leq k \leq n$, using the shorthand notations $(d)_l^m$ for the sequence d_l, \dots, d_m and s_1^k for a sequence of s_1 of length k , we define the following operations:

$$\begin{aligned} \mathbf{cons}_s &\in \mathcal{F}_{s_1^k, s} \\ \mathbf{enqueue}_s &\in \mathcal{F}_{s, s_1, s} \\ \mathbf{enqueue-at}_s &\in \mathcal{F}_{s, s_1, s', s} \end{aligned}$$

$$\begin{aligned}
\mathbf{push}_s &\in \mathcal{F}_{s,s_1,s} \\
\mathbf{push-at}_s &\in \mathcal{F}_{s,s_1,s',s} \\
\mathbf{remove}_s &\in \mathcal{F}_{s,s} \\
\mathbf{remove-at}_s &\in \mathcal{F}_{s,s',s} \\
\mathbf{peek}_s &\in \mathcal{F}_{s,s_1} \\
\mathbf{peek-at}_s &\in \mathcal{F}_{s,s',s_1} \\
\mathbf{free}_s &\in \mathcal{F}_{s,s'} \\
\mathbf{used}_s &\in \mathcal{F}_{s,s'} \\
\mathbf{cons}_s^{\mathcal{A}}(\langle (d)_1^k \rangle) &= \langle (d)_1^k \rangle \\
\mathbf{enqueue}_s^{\mathcal{A}}(\langle (d)_1^k \rangle, d') &= \begin{cases} \langle (d)_1^k, d' \rangle & \text{if } k < n \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{enqueue-at}_s^{\mathcal{A}}(\langle (d)_1^k \rangle, l, d') &= \begin{cases} \langle (d)_1^{k-l}, d', (d)_{k-l+1}^k \rangle & \text{if } 0 \leq l \leq k < n \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{push}_s^{\mathcal{A}}(\langle (d)_1^k \rangle, d') &= \begin{cases} \langle d', (d)_1^k \rangle & \text{if } k < n \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{push-at}_s^{\mathcal{A}}(\langle (d)_1^k \rangle, l, d') &= \begin{cases} \langle (d)_1^l, d', (d)_{l+1}^k \rangle & \text{if } 0 \leq l \leq k < n \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{remove}_s^{\mathcal{A}}(\langle (d)_1^k \rangle) &= \begin{cases} \langle (d)_2^k \rangle & \text{if } k \geq 1 \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{remove-at}_s^{\mathcal{A}}(\langle (d)_1^k \rangle, l) &= \begin{cases} \langle (d)_1^l, (d)_{l+2}^k \rangle & \text{if } 0 \leq l < k \geq 1 \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{peek}_s^{\mathcal{A}}(\langle (d)_1^k \rangle) &= \begin{cases} d_1 & \text{if } k \geq 1 \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{peek-at}_s^{\mathcal{A}}(\langle (d)_1^k \rangle, l) &= \begin{cases} d_{l+1} & \text{if } 0 \leq l < k \geq 1 \\ \varepsilon & \text{otherwise} \end{cases} \\
\mathbf{free}_s^{\mathcal{A}}(\langle (d)_1^k \rangle) &= n - k \\
\mathbf{used}_s^{\mathcal{A}}(\langle (d)_1^k \rangle) &= k.
\end{aligned}$$

The non-orthodox operations that allow the buffer to be accessed in the middle will be needed for an efficient transformation of some language constructs discussed in Section 9.2.2.

Tagged Union

For the tagged union defined in Section 5.3.4, we define three operations. Let $s_1, \dots, s_n \in \mathcal{S}_\beta$ for some $n > 0$, and let the corresponding carriers be $\mathcal{D}_{s_1}^{\mathcal{A}}, \dots, \mathcal{D}_{s_n}^{\mathcal{A}}$. Furthermore, let the sorts $s, b \in \mathcal{S}_\beta$ have the carriers $\mathcal{D}_s^{\mathcal{A}} = \mathbb{U}_{\mathcal{D}_{s_1}^{\mathcal{A}} \dots \mathcal{D}_{s_n}^{\mathcal{A}}}$ and $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$. We define the following operations for all $1 \leq k \leq n$:

$$\begin{aligned}
\mathbf{cons}_{s,k} &\in \mathcal{F}_{s_k,s} \\
\mathbf{component}_{s,k} &\in \mathcal{F}_{s,s_k}
\end{aligned}$$

$$\begin{aligned}
& \mathbf{defined}_{s,k} \in \mathcal{F}_{s,b} \\
& \mathbf{cons}_{s,k}^{\mathcal{A}}(d) = \langle k, d \rangle \\
& \mathbf{component}_{s,k}^{\mathcal{A}}(\langle l, d_l \rangle) = \begin{cases} d_l & \text{if } l = k \\ \varepsilon & \text{otherwise} \end{cases} \\
& \mathbf{defined}_{s,k}^{\mathcal{A}}(\langle l, d_l \rangle) = \begin{cases} \top & \text{if } l = k \\ \perp & \text{otherwise} \end{cases}
\end{aligned}$$

6.3.6 Type Conversions

We define three kinds of type conversion operations. The first kind will convert values between simple elementary types. The other two kinds are specific to data type systems having constraints. It is possible to convert values between types that differ only in the constraint, or between structured types whose component types can be pairwise converted to each other. The implementation in [41] performs some static analysis and does not allow conversions between types that have no common values.

The formal definition for the generic conversions is simple: For all $s, s' \in \mathcal{S}_\beta$ with the carriers $\mathcal{D}_s^{\mathcal{A}}$ and $\mathcal{D}_{s'}^{\mathcal{A}}$, we define

$$\begin{aligned}
& \mathbf{convert}_{s,s'} \in \mathcal{F}_{s,s'} \\
& \mathbf{convert}_{s,s'}^{\mathcal{A}}(d) = \begin{cases} d & \text{if } d \in \mathcal{D}_{s'}^{\mathcal{A}} \\ \langle k, d \rangle & \text{if } \mathcal{D}_{s'}^{\mathcal{A}} = \cup_{\mathcal{D}_{s_1}^{\mathcal{A}} \dots \mathcal{D}_{s_n}^{\mathcal{A}}} \text{ and } d \in \mathcal{D}_{s_k}^{\mathcal{A}} \text{ for a unique } k \\ d' & \text{if } \mathcal{D}_s^{\mathcal{A}} = \cup_{\mathcal{D}_{s_1}^{\mathcal{A}} \dots \mathcal{D}_{s_n}^{\mathcal{A}}}, d = \langle k, d' \rangle \text{ and } d' \in \mathcal{D}_{s'}^{\mathcal{A}} \\ \varepsilon & \text{otherwise} \end{cases}
\end{aligned}$$

We use the short-hand notation $(d)_1^n$ for the sequence d_1, \dots, d_n . For tuple sorts $s, s' \in \mathcal{S}_\beta$ of the respective component sorts $s_1, \dots, s_n \in \mathcal{S}_\beta$ and $s'_1, \dots, s'_n \in \mathcal{S}_\beta$, with

$$\mathcal{D}_s^{\mathcal{A}} = \mathbb{T}_{\mathcal{D}_{s_1}^{\mathcal{A}} \dots \mathcal{D}_{s_n}^{\mathcal{A}}} \quad \text{and} \quad \mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{T}_{\mathcal{D}_{s'_1}^{\mathcal{A}} \dots \mathcal{D}_{s'_n}^{\mathcal{A}}},$$

we extend the type conversion operation with the rule

$$\mathbf{convert}_{s,s'}^{\mathcal{A}}(\langle (d)_1^n \rangle) = \begin{cases} \mathbf{cons}_{s'}^{\mathcal{A}}(\langle (d')_1^n \rangle) & \text{if } d'_k = \mathbf{convert}_{s_k, s'_k}^{\mathcal{A}}(d_k) \neq \varepsilon \\ \varepsilon & \text{otherwise,} \end{cases}$$

quantifying the condition for all $1 \leq k \leq n$.

Similarly, for array sorts $s, s' \in \mathcal{S}_\beta$ with the index sort $s_x \in \mathcal{S}_\beta$ and the item sorts $s_e, s'_e \in \mathcal{S}_\beta$ and with the carriers such that

$$\mathcal{D}_s^{\mathcal{A}} = \mathbb{A}_{\mathcal{D}_{s_x}^{\mathcal{A}}, \mathcal{D}_{s_e}^{\mathcal{A}}} \quad \text{and} \quad \mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{A}_{\mathcal{D}_{s_x'}^{\mathcal{A}}, \mathcal{D}_{s'_e}^{\mathcal{A}}} \quad \text{and} \quad |\mathcal{D}_{s_x}^{\mathcal{A}}| = n,$$

we define

$$\mathbf{convert}_{s,s'}^{\mathcal{A}}(\langle (d)_1^n \rangle) = \begin{cases} \mathbf{cons}_{s'}^{\mathcal{A}}(\langle (d')_1^n \rangle) & \text{if } d'_k = \mathbf{convert}_{s_e, s'_e}^{\mathcal{A}}(d_k) \neq \varepsilon \\ \varepsilon & \text{otherwise.} \end{cases}$$

When $\mathcal{D}_s^{\mathcal{A}}$ and $\mathcal{D}_{s'}^{\mathcal{A}}$ are one of the simple types \mathbb{B} , \mathbb{K} , \mathbb{I} or \mathbb{E}_N defined in Section 5.2, conversions are based on numeric interpretations of the values. Let $s'' \in \mathcal{S}_\beta$ with $\mathcal{D}_{s''}^{\mathcal{A}} = \mathbb{I}$.

$$\begin{aligned} \mathbf{int}_s &\in \mathcal{F}_{s,s''} \\ \mathbf{int}_s^{\mathcal{A}}(d) &= \begin{cases} 0 & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{B} \text{ and } d = \perp \\ 1 & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{B} \text{ and } d = \top \\ c(d) & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{K} \\ d & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{I} \\ n_i & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{E}_N \text{ and } d = \langle i, n_i \rangle \\ \varepsilon & \text{otherwise} \end{cases} \\ \mathbf{ext}_s^{\mathcal{A}} &\in \mathcal{F}_{s'',s} \\ \mathbf{ext}_s^{\mathcal{A}}(d) &= \begin{cases} \perp & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{B} \text{ and } d = 0 \\ \top & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{B} \text{ and } d = 1 \\ d' & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{K} \text{ and } d = c(d') \\ d & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{I} \\ \langle i, n_i \rangle & \text{if } \mathcal{D}_s^{\mathcal{A}} = \mathbb{E}_N \text{ and } \exists i \in N : d = n_i \\ \varepsilon & \text{otherwise} \end{cases} \end{aligned}$$

The implementation in [41] combines the two operations as

$$\mathbf{ext}_s^{\mathcal{A}}(\mathbf{int}_{s'}^{\mathcal{A}}) : \mathcal{D}_{s'}^{\mathcal{A}} \rightarrow \check{\mathcal{D}}_s^{\mathcal{A}}$$

and provides uniform syntax for all type conversion operations.

6.4 Operations on Multi-Set Sorts

All operations introduced so far have basic sorts as their carriers. As the arc expressions in Algebraic System Nets are sorted over sets of multi-sets, we will need to define some operations with multi-set sorts as their carriers.

6.4.1 Multi-Set Constructor

For $s, s' \in \mathcal{S}_\beta$ with the carriers $\mathcal{D}_s^{\mathcal{A}}$ and $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{I}$, respectively, the multi-set constructor

$$\begin{aligned} \mathbf{mset}_s &\in \mathcal{F}_{s,s',\mu(s)} \\ \mathbf{mset}_s^{\mathcal{A}}(d, n) &= \begin{cases} \{\langle d, n \rangle\} \cup \bigcup_{d' \in \mathcal{D}_s^{\mathcal{A}} \setminus \{d\}} \{\langle d', 0 \rangle\} & \text{if } d \neq \varepsilon \text{ and } n \geq 0 \\ \varepsilon & \text{otherwise} \end{cases} \end{aligned}$$

constructs a multi-set containing some number of a single element.

6.4.2 Empty Multi-Set

For $s \in \mathcal{S}_\beta$ with a carrier $\mathcal{D}_s^{\mathcal{A}}$, the empty multi-set over s is defined as

$$\begin{aligned} \mathbf{empty}_s &\in \mathcal{F}_{\lambda, \mu(s)} \\ \mathbf{empty}_s^{\mathcal{A}}() &= \left\{ \langle d, 0 \rangle \parallel d \in \mathcal{D}_s^{\mathcal{A}} \right\}. \end{aligned}$$

6.4.3 Multi-Set Sum and Filter

For $s, s', b \in \mathcal{S}_\beta$ with the respective carriers $\mathcal{D}_s^{\mathcal{A}}$, $\mathcal{D}_{s'}^{\mathcal{A}}$ and $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$, and for the terms $T \in \mathbf{T}_b^{\mathbf{S}^\mu}(\mathcal{V})$ and $U \in \mathbf{T}_{\mu(s')}^{\mathbf{S}^\mu}(\mathcal{V})$ where $\mathcal{V} = \mathcal{V}_{s_1} \cup \dots \cup \mathcal{V}_{s_n} \cup \mathcal{V}_s$ for some $n \geq 0$ such that $x \in \mathcal{V}_s$, $x_i \in \mathcal{V}_{s_i}$ for $s_i \in \mathcal{S}$ and $1 \leq i \leq n$, we define the following operations:

$$\begin{aligned} \mathbf{sum}_{x,T,U} &\in \mathcal{F}_{s_1, \dots, s_n, \mu(s')} \\ \mathbf{filter}_{x,T} &\in \mathcal{F}_{\mu(s), s_1, \dots, s_n, \mu(s)} \\ \mathbf{sum}_{x,T,U}^{\mathcal{A}}(d_1, \dots, d_n) &= \begin{cases} \varepsilon & \text{if } g(V_{\mathcal{D}_s^{\mathcal{A}}}^{\mathcal{A}}(\mathcal{V}), T) \\ \varepsilon & \text{if } g(V_{\mathcal{D}_{s'}^{\mathcal{A}}}^{\mathcal{A}}(\mathcal{V}), U) \\ \mathbf{empty}_{s'}^{\mathcal{A}}() & \text{if } V_{\mathcal{D}_s^{\mathcal{A}}, T}^{\mathcal{A}}(\mathcal{V}) = \emptyset \\ \sum_{v \in V_{\mathcal{D}_s^{\mathcal{A}}, T}^{\mathcal{A}}(\mathcal{V})} e_v^{\mathcal{A}}(U) & \text{otherwise} \end{cases} \\ \mathbf{filter}_{x,T}^{\mathcal{A}}(\mu, d_1, \dots, d_n) &= \begin{cases} \varepsilon & \text{if } g(V_\mu^{\mathcal{A}}(\mathcal{V}), T) \\ \mathbf{empty}_s^{\mathcal{A}}() & \text{if } V_\mu^{\mathcal{A}}(\mathcal{V}) = \emptyset \\ \sum_{v \in V_\mu^{\mathcal{A}}(\mathcal{V})} \mathbf{mset}_s^{\mathcal{A}}(v(x), \mu_v(v(x), T)) & \text{otherwise} \end{cases} \end{aligned}$$

where the summations are multi-set unions as in Definition 4.10 and

$$\begin{aligned} V_{\mathcal{D}_s^{\mathcal{A}}}^{\mathcal{A}}(\mathcal{V}) &= \bigcup_{d \in \mathcal{D}_s^{\mathcal{A}}} \left\{ \{ \langle x_1, d_1 \rangle, \dots, \langle x_n, d_n \rangle, \langle x, d \rangle \} \right\} \\ V_{\mathcal{D}_s^{\mathcal{A}}, T}^{\mathcal{A}}(\mathcal{V}) &= \left\{ v \in V_{\mathcal{D}_s^{\mathcal{A}}}^{\mathcal{A}}(\mathcal{V}) \parallel e_v^{\mathcal{A}}(T) = \top \right\} \\ V_\mu^{\mathcal{A}}(\mathcal{V}) &= \bigcup_{d \in \mu} \left\{ \{ \langle x_1, d_1 \rangle, \dots, \langle x_n, d_n \rangle, \langle x, d \rangle \} \right\} \\ g(V, T) &\equiv \exists v \in V : e_v^{\mathcal{A}}(T) = \varepsilon \\ \mu_v(d, T) &= \begin{cases} \mu(d) & \text{if } e_v^{\mathcal{A}}(T) = \top \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The sum and the filter operations overlap in functionality. In typical models, the multi-set sum is used mostly for constructing initial markings, whereas the filter is more common in temporal logic formulae and on output arc expressions with the μ not being a function of a multi-set sum.

6.4.4 Multi-Set Transformations

Item Mapping

For $s \in \mathcal{S}_\beta$ with the carrier $\mathcal{D}_s^{\mathcal{A}}$, for $s' \in \mathcal{S}_\mu$ and for the term $T \in \mathbf{T}_s^{\mathcal{S}\mu}(\mathcal{V})$ with $\mathcal{V} = \mathcal{V}_{s_1} \cup \dots \cup \mathcal{V}_{s_n} \cup \mathcal{V}_s$ for some $n \geq 0$ such that $x \in \mathcal{V}_s$, $x_i \in \mathcal{V}_{s_i}$ for $s_i \in \mathcal{S}$ and $1 \leq i \leq n$, we define the multi-set item mapping operation

$$\mathbf{map}_{x,T} \in \mathcal{F}_{s',s_1,\dots,s_n,\mu(s)}$$

$$\mathbf{map}_{x,T}^{\mathcal{A}}(\mu, d_1, \dots, d_n) = \begin{cases} \varepsilon & \text{if } g(V_\mu^{\mathcal{A}}(\mathcal{V}), T) \\ \varepsilon & \text{if } \exists d \in \mu : \mu(d) \notin \mathbb{I} \\ \mathbf{empty}_s^{\mathcal{A}}() & \text{if } V_\mu^{\mathcal{A}}(\mathcal{V}) = \emptyset \\ \sum_{v \in V_\mu^{\mathcal{A}}(\mathcal{V})} \mathbf{mset}_s^{\mathcal{A}}(e_v^{\mathcal{A}}(T), \mu(v(x))) & \text{otherwise} \end{cases}$$

where the summation over v is a multi-set union as in Definition 4.10 and

$$V_\mu^{\mathcal{A}}(\mathcal{V}) = \bigcup_{d \in \mu} \{ \langle x_1, d_1 \rangle, \dots, \langle x_n, d_n \rangle, \langle x, d \rangle \}$$

$$g(V, T) \equiv \exists v \in V : e_v^{\mathcal{A}}(T) = \varepsilon.$$

Multiplicity Mapping

For $s, s' \in \mathcal{S}_\beta$ with the carriers $\mathcal{D}_s^{\mathcal{A}}$ and $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{I}$ and for the term $T \in \mathbf{T}_s^{\mathcal{S}\mu}(\mathcal{V})$ with $\mathcal{V} = \mathcal{V}_{s_1} \cup \dots \cup \mathcal{V}_{s_n} \cup \mathcal{V}_s \cup \mathcal{V}_{s'}$ for some $n \geq 0$ such that $x \in \mathcal{V}_s$, $y \in \mathcal{V}_{s'}$ and $x_i \in \mathcal{V}_{s_i}$ for $s_i \in \mathcal{S}$ and $1 \leq i \leq n$, we define the multi-set multiplicity mapping operation

$$\mathbf{mmap}_{x,y,T} \in \mathcal{F}_{\mu(s),s_1,\dots,s_n,\mu(s)}$$

$$\mathbf{mmap}_{x,y,T}^{\mathcal{A}}(\mu, d_1, \dots, d_n) = \begin{cases} \varepsilon & \text{if } h(V_\mu^{\mathcal{A}}(\mathcal{V}), T) \\ \varepsilon & \text{if } \exists d \in \mu : \mu(d) \notin \mathbb{I} \\ \mathbf{empty}_s^{\mathcal{A}}() & \text{if } V_\mu^{\mathcal{A}}(\mathcal{V}) = \emptyset \\ \sum_{v \in V_\mu^{\mathcal{A}}(\mathcal{V})} \mathbf{mset}_s^{\mathcal{A}}(v(x), e_v^{\mathcal{A}}(T)) & \text{otherwise} \end{cases}$$

where the summation over v is a multi-set union as in Definition 4.10 and

$$V_\mu^{\mathcal{A}}(\mathcal{V}) = \bigcup_{d \in \mu} \{ \langle x_1, d_1 \rangle, \dots, \langle x_n, d_n \rangle, \langle x, d \rangle, \langle y, \mu(d) \rangle \}$$

$$h(V, T) \equiv \exists v \in V : e_v^{\mathcal{A}}(T) = \varepsilon \vee e_v^{\mathcal{A}}(T) < 0.$$

6.4.5 Union and Intersection

For $s \in \mathcal{S}_\beta$, we define the following multi-set operations:

$$\mathbf{union}_s \in \mathcal{F}_{\mu(s),\mu(s),\mu(s)}$$

$$\mathbf{minus}_s \in \mathcal{F}_{\mu(s),\mu(s),\mu(s)}$$

$$\mathbf{union}_s^{\mathcal{A}}(\mu_1, \mu_2) = \mu_1 + \mu_2$$

$$\mathbf{minus}_s^{\mathcal{A}}(\mu_1, \mu_2) = \mu_1 - \mu_2$$

6.4.6 Scalar Multiplication

With the operations defined so far, it is possible to construct all the multi-sets needed in practice. The following operation will not add any expressive power to the set of defined operations. Multi-set-valued terms can be shortened by making use of this operation, just like scalar expressions can be shortened by making use of grouping, e.g.: $a \cdot b + a \cdot c = a \cdot (b + c)$.

For $s \in \mathcal{S}_\mu$ and $s' \in \mathcal{S}_\beta$ with the carrier $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{I}$, we define the scalar multiplication of a multi-set by an integer as follows:

$$\begin{aligned} \mathbf{mul}_s &\in \mathcal{F}_{s,s',s} \\ \mathbf{mul}_s^{\mathcal{A}}(\mu, d) &= \begin{cases} d \cdot \mu & \text{if } d \geq 0 \\ \varepsilon & \text{otherwise} \end{cases} \end{aligned}$$

6.4.7 Comparison

For $s \in \mathcal{S}_\mu$ and $b \in \mathcal{S}_\beta$ with the carriers $\mathcal{D}_s^{\mathcal{A}}$ and $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$, we define the following binary relations $\mathcal{D}_s^{\mathcal{A}} \times \mathcal{D}_s^{\mathcal{A}} \rightarrow \mathbb{B}$ for testing multi-set equality and containment:

$$\begin{aligned} \mathbf{equalset}_s &\in \mathcal{F}_{s,s,b} \\ \mathbf{subset}_s &\in \mathcal{F}_{s,s,b} \\ \mathbf{equalset}_s^{\mathcal{A}}(\mu_1, \mu_2) &= \begin{cases} \top & \text{if } \mu_1 = \mu_2 \\ \perp & \text{otherwise} \end{cases} \\ \mathbf{subset}_s^{\mathcal{A}}(\mu_1, \mu_2) &= \begin{cases} \top & \text{if } \mu_1 \leq \mu_2 \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

6.4.8 Cardinality

For $s, s' \in \mathcal{S}_\beta$ with the carriers $\mathcal{D}_s^{\mathcal{A}}$ and $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{I}$, we define the following operation:

$$\begin{aligned} \mathbf{card}_s &\in \mathcal{F}_{\mu(s),s'} \\ \mathbf{card}_s^{\mathcal{A}}(\mu) &= \begin{cases} |\mu| & \text{if } |\mu| \in \mathbb{I} \\ \varepsilon & \text{otherwise} \end{cases} \end{aligned}$$

6.5 Short-Circuit Operations

Some programming languages make use of a technique called *short-circuit evaluation*. Under some circumstances, it suffices to evaluate only part of an expression. The built-in binary operators for logical conjunction `&&` and disjunction `||` and the ternary if-then-else operator `?:` of the programming languages C [28] and C++ [27] are a good example. The binary operators can be represented with the ternary operator:

$$\begin{aligned} a \&\&b &= a ? b : 0 \\ a \|\| b &= a ? 1 : b \end{aligned}$$

When the ternary operation is evaluated, its leftmost operand will be evaluated first. If it evaluates to zero or false, the result of the ternary operation will be the evaluation of its third operand; otherwise the outcome will be the evaluation of the second operand.

In [41], there are short-circuit versions of the logical conjunction and disjunction operators and a generalised $(n + 1)$ -ary variant of the selection operator $?$;, which takes a left-hand-side term with a carrier of n elements, and n right-hand-side terms separated by colons, one of which will be selected for evaluation based on the outcome of the evaluation of the left-hand-side argument. The selected right-hand-side term will determine the outcome of the operation, as defined in Definition 4.7. We shall define only the selection operation here, since short-circuit disjunctions and conjunctions can be directly represented with it.

For $s' \in S_\beta$ and $s \in S$ with the carriers $\mathcal{D}_{s'}^{\mathcal{A}}$ and $\mathcal{D}_s^{\mathcal{A}}$ and $n = |\mathcal{D}_{s'}^{\mathcal{A}}|$, we define

$$\mathbf{select}_{s',s^n} \in \mathcal{G}_{s',s^n}$$

where s^n denotes the sequence $\underbrace{s, \dots, s}_{n \text{ times}}$.

Part III

Implementing and Using the Analyser

Chapter 7

Implementing the Analyser

There are some implementation details we would like to point out, since they affect the exact semantics of the class of algebraic system nets we have implemented in [41]. Moreover, the way how an algorithm is implemented can be important or even crucial in practice, even if the difference is just an n -fold increase in execution speed or decrease in memory usage for some constant n . We will discuss some optimisations that have turned out to be extremely useful in practice.

7.1 Transition Instance Analysis

Definition 4.17, the transition pre-enabling rule, does not mention how the assignments $v \in V^{\mathcal{A}}(\mathcal{V})$ that pre-enable a transition $t \in \mathcal{T}$ in a marking M can be found. A naïve approach would be to enumerate through all assignments, without paying any attention to the marking M . Doing so is possible when all input variables of a transition have a finite carrier, but such an algorithm would have exponential time complexity.

We will give some definitions before we shed some light on the algorithm implemented in [41]. There are some pathological cases where the algorithm fails to find values for all variables needed for evaluating the transition expressions, but the algorithm will report an error in these cases.

Our algorithm makes use of three definitions. The first of them defines *assignment candidates*. Given an algebraic term and a value the term should evaluate to, the definition yields a set of variable bindings under which the algebraic term could be *compatible* with the value, i.e. evaluate either to the value or fail to be evaluated because of an undefined variable.

Our definition of assignment candidates reflects our implementation [41]. Although the definition could cover all reversible algebraic operations—operations computable with injective mappings—it only covers variables and constructors for structured data types, defined in Section 6.3.5. For instance, the successor and predecessor operations defined in Section 6.3.2 are outside its scope.¹

¹This choice not only benefits the lazy programmer but also persons studying a model, who do not need to know which operations are considered injective or reversible.

Definition 7.1 (Assignment Candidates) Let $\mathbf{S}_\mu = \langle \mathcal{S}_\beta \cup \mathcal{S}_\mu, \mathcal{F}, \mathcal{G}, \mu \rangle$ be a multi-set signature and \mathcal{A} be an \mathbf{S}_μ -algebra. Let $T \in \mathbf{T}_s^{\mathbf{S}_\mu}(\mathcal{V})$ be a term, $v \in V^{\mathcal{A}}(\mathcal{V})$ a valuation and $d \in \mathcal{D}_s^{\mathcal{A}}$ a value. The set of assignment candidates for T , v and d , denoted

$$A_v(T, d),$$

is the minimal subset of $\bigcup_{s \in \mathcal{S}_\beta} (\mathcal{V}_s \rightarrow \mathcal{D}_s^{\mathcal{A}})$ defined inductively by the following rules for each variable $x \in \mathcal{V}_s$ such that $v(x) = \varepsilon$:

1. if $T = x$, then $\langle x, d \rangle \in A_v(T, d)$
2. if for some $n \geq 0$, $T = \mathbf{cons}_s(T_1, \dots, T_n)$ and $d = \langle d_1, \dots, d_n \rangle$, then for each $1 \leq k \leq n$, $A_v(T_k, d_k) \subseteq A_v(T, d)$.

Definition 7.2 (Value Compatibility) Let $\mathbf{S}_\mu = \langle \mathcal{S}_\beta \cup \mathcal{S}_\mu, \mathcal{F}, \mathcal{G}, \mu \rangle$ be a multi-set signature and \mathcal{A} be an \mathbf{S}_μ -algebra. Let $s \in \mathcal{S}_\beta$ be a sort with the carrier $\mathcal{D}_s^{\mathcal{A}}$, $T \in \mathbf{T}_s^{\mathbf{S}_\mu}(\mathcal{V})$ a term and $v \in V^{\mathcal{A}}(\mathcal{V})$ a valuation. A value $d \in \mathcal{D}_s^{\mathcal{A}}$ is compatible with T under v , denoted

$$T \approx_v d,$$

if at least one of the following conditions holds:

1. $e_v^{\mathcal{A}}(T) = d$
2. $T = \mathbf{convert}_{s',s}(T')$ and $T' \approx_v d$
3. for some $n \geq 0$, $T = \mathbf{cons}_s(T_1, \dots, T_n)$, $d = \langle d_1, \dots, d_n \rangle$, and for each $1 \leq k \leq n$, either $T_k \approx_v d_k$ or $e_v^{\mathcal{A}}(T_k) = \varepsilon$.

Definition 7.3 (Token Compatibility) Let $\mathbf{S}_\mu = \langle \mathcal{S}_\beta \cup \mathcal{S}_\mu, \mathcal{F}, \mathcal{G}, \mu \rangle$ be a multi-set signature and \mathcal{A} be an \mathbf{S}_μ -algebra. Let $s \in \mathcal{S}_\beta$ be a sort with the carrier $\mathcal{D}_s^{\mathcal{A}}$, $T \in \mathbf{T}_{\mu(s)}^{\mathbf{S}_\mu}(\mathcal{V})$ a term, and let $v \in V^{\mathcal{A}}(\mathcal{V})$ be an assignment. A multi-set $\mu \in \mathcal{M}(\mathcal{D}_s^{\mathcal{A}})$ is compatible with T under v , denoted

$$T \approx_v \mu,$$

if at least one of the following conditions holds:

1. $e_v^{\mathcal{A}}(T) = \mu$
2. $T = \mathbf{mset}_s(U, V)$ for some $U \in \mathbf{T}_s^{\mathbf{S}_\mu}(\mathcal{V})$ and $V \in \mathbf{T}_{s'}^{\mathbf{S}_\mu}(\mathcal{V})$ such that $n = e_v^{\mathcal{A}}(V) \geq 0$, and either $n = 0$ or there is $d \in \mu$ such that $U \approx_v d$ and $n = \mu(d)$ and $\mu(d') = 0$ for all $d' \in \mathcal{D}_s^{\mathcal{A}} \setminus \{d\}$.

Compute all pre-enabling assignments for transition t in marking M

```

V ← ∅                                assignments found so far
Vε ← ∅                               erroneous assignments found so far
S ← ∅                                unification stack
v ← ⋃x∈V{⟨x, ε⟩}                    current assignment
T ← NEXTTOKEN(M, t, v)                get next symbolic token, or NIL
do
  if T ≠ NIL
    do
      do
        v ← UNDEFINE(T, v)             undo EXTEND(T, μ, v)
        μ ← NEXT(M, T, v)              get next μ ≤ M : T ≈v μ ∧ evA(i(t)) ≠ ⊥
        if μ = NIL goto back          backtrack if no such μ found
        v ← EXTEND(T, μ, v)            extend v with Av(T, d) for d ∈ μ
      until T ≈v μ                    retry if incompatible
      and ∀⟨T', μ'⟩ ∈ S : T' ≈v μ'   retry if stack incompatible
      and evA(i(t)) ≠ ⊥              retry if guard disabled
      RESERVE(M, μ)                    reserve μ from M: M ← M - μ
      PUSH(S, ⟨T, μ⟩)                  store the (partially) unified tokens
      T ← NEXTTOKEN(M, t, v)           get next symbolic token, or NIL
    until T = NIL                    loop until all symbolic tokens unified
  if S = ∅
    goto back                          end of search; clean up
  if ∃⟨T', μ'⟩ ∈ S : evA(T') ≠ μ'
  or ∃⟨p, t⟩ ∈ F : evA(i(⟨p, t⟩)) = ε
  or evA(i(t)) ≠ ⊤
    Vε ← Vε ∪ {v}                    store the erroneous assignment
  else
    V ← V ∪ {v}                          store the completed assignment
back:
  if T ≠ NIL
    RELEASETOKEN(T)                      make T available for NEXTTOKEN
  if S ≠ ∅
    ⟨T, μ⟩ ← POP(S)                       backtrack
    RELEASE(M, μ)                          put μ back to M: M ← M + μ
until S = ∅
return ⟨V, Vε⟩                       search completed

```

Figure 7.1: Transition Instance Analysis

7.1.1 Rough Description of the Algorithm

Our transition instance analysis algorithm, presented in a simplified form in Figure 7.1, starts with a transition $t \in \mathcal{T}$, a marking M , an initially empty assignment $\nu : \mathcal{V} \rightarrow \{\varepsilon\}$ and an initially empty unification stack of pairs of multi-sets and compatible *symbolic tokens*, terms of the form $\mathbf{mset}_s(U, \mathbf{constant}_n())$.

Each call to NEXTTOKEN iterates through all expressions $i(f)$ on the input arcs $f = \langle p, t \rangle \in \mathcal{F}$, trying to convert them to symbolic tokens. Those terms whose multiplicity cannot be evaluated are skipped in the iteration. NEXTTOKEN keeps track on the symbolic tokens it has returned, and it will not return the same token again until it has been released by invoking RELEASETOKEN.

Each symbolic token found by the algorithm will be pushed on the unification stack together with a speculative token, a compatible token that will be reserved (subtracted) from M . All assignment candidates in the symbolic tokens are set according to the speculative token.

The iteration over the input arc expressions continues until all symbolic tokens have been extracted. At that point, the assignment is either pre-enabling or erroneous. Our implementation augments the set of erroneous assignments also whenever a term evaluates to ε due to something else than an undefined variable.

Upon reaching the bottom of this depth-first search, the algorithm backtracks and changes the speculative token μ assigned to the symbolic token T by invoking NEXT. Once all possible bindings for T have been covered, the token is thrown away, and another pair $\langle T, \mu \rangle$ is pulled out of the unification stack. When the unification stack runs out of tokens, all possibilities—pre-enabling or erroneous assignments—have been discovered, and the algorithm terminates.

The algorithm could take more advantage of the structure of the symbolic tokens and process the arcs in a certain order, minimising the total number of speculative token bindings required. Currently the algorithm only utilises the transition guard $i(t)$: when a conjunctive term of the transition guard evaluates to \perp , the algorithm can prune the branch of the search tree and bind the next speculative token to the symbolic token on top of the unification stack.

In spite of its limitations, the algorithm finds all pre-enabling assignments except in a few pathological cases, e.g. when a transition has only one input arc

$$\mathbf{union}_s(\mathbf{mset}_s(x, y), \mathbf{mset}_s(y, x))$$

where $x, y \in \mathcal{V}_s$ and $\mathcal{D}_s^{\mathbb{A}} = \mathbb{I}$ (the integer type defined in Section 5.2.3). Since there are no assignment candidates for x and y , no symbolic tokens can be extracted from the input arc.

An Example

We shall illustrate the algorithm with a simple example, presented in Figure 7.2. For the sake of simplicity, place inscriptions are omitted from the example. For a marking M with

$$\begin{aligned} M(\mathbf{A}) &= \{\langle 0, 0 \rangle, \langle 1, 1 \rangle, \langle 2, 1 \rangle, \langle 3, 1 \rangle\} \\ M(\mathbf{B}) &= \{\langle 0, 3 \rangle, \langle 1, 2 \rangle, \langle 2, 1 \rangle, \langle 3, 0 \rangle\} \end{aligned}$$

$$\begin{aligned}
\Sigma &= \langle \mathcal{N}, \mathcal{A}, \mathcal{V}, i \rangle \\
\mathcal{N} &= \langle \mathcal{P}, \mathcal{T}, \mathcal{F} \rangle \\
s &\in \mathcal{S}_\beta \\
\mathcal{D}_s^{\mathcal{A}} &= \{0, 1, 2, 3\} \\
\mathcal{P}_{\mu(s)} &= \{\mathbf{A}, \mathbf{B}\} \\
\mathcal{P} &= \mathcal{P}_{\mu(s)} \\
\mathcal{T} &= \{\mathcal{T}\} \\
\mathcal{F} &= \{\langle \mathbf{A}, \mathcal{T} \rangle, \langle \mathbf{B}, \mathcal{T} \rangle\} \\
\mathcal{V} &= \mathcal{V}'_s \\
\mathcal{V}'_s &= \{x, y\} \\
i(\langle \mathbf{A}, \mathcal{T} \rangle) &= \mathbf{union}_s(\mathbf{mset}_s(x, \mathbf{constant}_1()), \mathbf{mset}_s(\mathbf{succ}(x), \mathbf{constant}_1())) \\
i(\langle \mathbf{B}, \mathcal{T} \rangle) &= \mathbf{mset}_s(y, \mathbf{pred}(x)) \\
i(\mathcal{T}) &= \mathbf{constant}_\top()
\end{aligned}$$

Figure 7.2: An Example Model for Transition Instance Analysis

the instance analysis for transition \mathcal{T} proceeds as follows.

The first call to NEXTTOKEN yields the term $\mathbf{mset}_s(x, \mathbf{constant}_1())$ from the arc $\langle \mathbf{A}, \mathcal{T} \rangle$. Calling UNDEFINE has no effect on the first round of the main loop. The subroutine NEXT iterates through the marking of the place \mathbf{A} . On the first call, it yields $\{\langle 1, 1 \rangle\}$, one token of the value 1. The assignment is extended by setting $v = \{x, 1\}$. The compatibility checks pass for the symbolic token under the augmented assignment, and the first token is fully unified. It is reserved from the marking and pushed on the unification stack.

The call to NEXTTOKEN yields the term $\mathbf{mset}_s(\mathbf{succ}(x), \mathbf{constant}_1())$ from the same arc $\langle \mathbf{A}, \mathcal{T} \rangle$. This term evaluates to $\langle 2, 1 \rangle$ under the current assignment. Calling UNDEFINE has no effect, since this symbolic token has not been processed yet by the algorithm. Now the subroutine NEXT has two tokens to choose from, $\langle 2, 1 \rangle$ and $\langle 3, 1 \rangle$, since $\langle 1, 1 \rangle$ was reserved from $M(\mathbf{A})$ for the first symbolic token. The token $\langle 2, 1 \rangle$ matches, and the tokens in the unification stack remain compatible under the assignment, since the assignment was not changed. Also the second token is reserved from the marking, and the tokens are pushed on the unification stack.

The third call to NEXTTOKEN yields the term $\mathbf{mset}_s(y, \mathbf{pred}(x))$ from the arc $\langle \mathbf{B}, \mathcal{T} \rangle$. As $\mathbf{pred}(x)$ evaluates to 0 under the current assignment, the call to NEXT has to pick a token of multiplicity 0, i.e. nothing, from the place \mathbf{B} . The assignment will not be augmented, since μ is the empty multi-set. The term and the tokens in the unification stack are compatible under this assignment, and also the transition guard $i(\mathcal{T})$ is satisfied. The token is reserved from the marking and the tokens are pushed on the unification stack.

Now, since there are no more unprocessed symbolic tokens, NEXTTOKEN

yields NIL and the second loop terminates. The completed assignment v is stored in the set of found assignments, and the last token and corresponding symbolic token are retrieved from the unification stack. Releasing the token to the marking has no effect, since it is empty. Likewise, the call to UNDEFINE has no effect, since the assignment was not augmented further from $v = \{\langle x, 1 \rangle\}$. The call to NEXT does not find any more token candidates: it yielded the only possibility (the empty token) already in the previous round. The algorithm backtracks.

At this step, the unification stack contains two tokens. The newer one, the one with $T = \mathbf{mset}_s(\mathbf{succ}(x), \mathbf{constant}_1())$, is pulled from the stack, and $\mu = \langle 2, 1 \rangle$ is released to the marking. The call to UNDEFINE has no effect on this term, since the assignment was not augmented when this symbolic token was unified. The call to NEXT does not find any more token candidates, and the algorithm backtracks further. Now the first symbolic token, $T = \mathbf{mset}_s(x, \mathbf{constant}_1())$, is pulled from the stack, and we are almost back to square one.

At this point, the call to UNDEFINE clears the assignment and NEXT results in $\mu = \langle 2, 1 \rangle$. Eventually the tokens are pushed on the stack, and the algorithm considers the next symbolic token, $\mathbf{mset}_s(\mathbf{succ}(x), \mathbf{constant}_1())$, and the only possible token, $\mu = \langle 3, 1 \rangle$. Also these will be pushed on the stack, and the algorithm will proceed to the third symbolic token, $\mathbf{mset}_s(y, \mathbf{pred}(x))$. Now $\mathbf{pred}(x)$ evaluates to 1, and the call to NEXT yields one of the three token kinds contained in the place **B**: $\langle 0, 1 \rangle$, $\langle 1, 1 \rangle$ and $\langle 2, 1 \rangle$. Let us assume the token is $\mu = \langle 0, 1 \rangle$. The call to EXTEND augments the assignment to $v = \{\langle x, 2 \rangle, \langle y, 0 \rangle\}$, and the tokens are pushed on the unification stack.

The second loop terminates, since there are no more unprocessed tokens. As all tokens in the unification stack are compatible with the current assignment v , the set of assignments is augmented with v . This procedure is repeated for all three tokens in **B**. Finally, the set of assignments is

$$V = \{\{\langle x, 1 \rangle\}, \{\langle x, 2 \rangle, \langle y, 0 \rangle\}, \{\langle x, 2 \rangle, \langle y, 1 \rangle\}, \{\langle x, 2 \rangle, \langle y, 2 \rangle\}\}$$

and the algorithm terminates after backtracking until the unification stack is empty.

This example shows that our unification algorithm does not necessarily provide values for all variables. If an output arc of the transition T needs to evaluate the variable y , the assignment $\{\langle x, 1 \rangle\}$ that pre-enables T will not enable T .

Our unification algorithm could probably be optimised considerably. The pattern matching algorithm presented in [45] would be advantageous in situations where the marking M does not change much between successive invocations of the algorithm. This could be useful when generating the set of reachable states in a depth-first order.

7.1.2 Unfolding to Place/Transition Nets

Some classes of high-level Petri Nets can be mapped to Place/Transition Nets [50], which consist of low-level places and transitions connected via weighted arcs. In a Place/Transition Net, places may contain a number of black tokens, which are moved around by the firing of low-level transitions. The transformation to a lower-

level formalism is motivated, since many reduction methods for reachability analysis are only defined for Place/Transition Nets.

High-level Petri Nets are usually *unfolded* to Place/Transition Nets by mapping each place $p \in \mathcal{P}_{\mu(s)}$ with $s \in \mathcal{S}_{\beta}$ to $|\mathcal{D}_s^{\mathcal{A}}|$ low-level places and each possible enabled transition instance $\langle t, \hat{v}, v \rangle$ with $t \in \mathcal{T}$ and $\hat{v}, v \in V^{\mathcal{A}}(\mathcal{V})$ and $e_{\hat{v}}^{\mathcal{A}}(i(t)) = \top$, to a low-level transition connected with the low-level places corresponding to the input and output effects $t_{\hat{v}}^{-}$ and t_v^{+} . See e.g. [37, Section 5.1] for an illustration of this approach.

It is worth noting that this kind of a transformation is independent of the initial marking and the set of reachable markings; the unfolded Place/Transition Net may contain many places whose markings will never change or transitions that will never fire. A more efficient unfolding would construct the set of low-level places incrementally, starting from the low-level places that are non-empty in the initial marking, and adding necessary low-level places for each transition instance. Also, some data types, such as the buffer data type, can be modelled with more efficient constructs in some special cases.

The algorithm described in the previous section can be fairly easily modified to perform unfolding. Instead of taking the speculative tokens from a given marking, the unfolding algorithm will consider an “infinite” marking. This has not been implemented yet.

7.2 The Expression Evaluator

Implementing an expression evaluator, an algorithm that computes $e_v^{\mathcal{A}}(T)$, is a straightforward task, once the interfaces have been defined. In an interpreter-based approach, it is customary to implement the expression evaluator in a set of functions that recursively call each other. An object-oriented implementation language, such as our choice C++, allows the expression evaluator to be implemented in virtual methods, which makes it fairly easy to add new types of subexpressions, each of which corresponds to an operation defined in Chapter 6.

The principle of an expression evaluator is very simple: it can be implemented as a procedure that takes an expression (an algebraic term) and an assignment (values for variables), and evaluates the expression by evaluating the subexpressions of each expression by calling itself recursively, and by using the evaluated values of the subexpressions for computing the value of the expression. The assignment is only needed for evaluating variable references.

7.2.1 Error Handling

What should happen when a subexpression is undefined, if for instance it contains a division by zero? One approach is to sweep the problem under the carpet by using special values for undefined subexpressions, such as the ε defined in our algebra, or the NaN (Not a Number) values used by many floating-point evaluators. Another way is to use an exception handling mechanism to abort the expression evaluator immediately when an error occurs.

Our implementation uses a special value together with a more specific error reporting mechanism. The expression evaluator is implemented as a virtual method that takes one parameter, an object containing the assignment for the variables. When an error occurs, the expression evaluator will return a special value and set an error code in the assignment object. The assignment object has a diagnostic method that will display the error code together with the complete assignment. This is a very useful way of reporting errors, especially those that occur during the transition instance analysis.

7.2.2 Optimisations

One cannot expect much performance from an interpreter-based expression evaluator. The performance can be improved by introducing optimisations. Some of them transform expressions to simpler equivalent ones, while others affect the expression evaluator itself.

Constant Folding

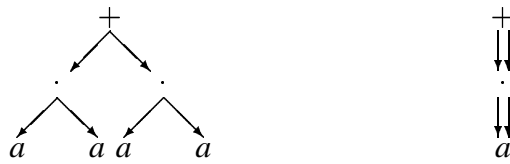
Evaluating a variable-free expression always yields the same result. If the expression is evaluated more than once, computing power will be wasted. Such expressions, called ground terms or constant expressions, can be evaluated by the expression parser and replaced with constants representing their values. This technique is called *constant folding*. [1]

Constant folding can also improve error checking. If there is an error in a constant expression, the error can be detected already in the parsing stage, and it is easy to associate a diagnostic message with the exact location of the error in the input. Would the error occur at a later stage, when evaluating the expression in a dynamic context, reporting it precisely would be more difficult. Also, if not all expressions will ever be evaluated, errors found during constant folding could go undetected.

Common Subexpression Elimination

Expressions can be viewed as trees, which are a special case of directed acyclic graphs. A well-known optimisation technique called common subexpression elimination [1] views the expression as a directed graph. Identical subexpressions are not represented by identical subtrees, but by multiple arcs leading to the same subgraph representing the subexpression. Instead of parsing an expression to a tree, it can be parsed to a directed acyclic graph. A simple example is illustrated in Figure 7.3.

Common subexpression elimination has many advantages, and a typical expression tends to have common subexpressions, at least variable references. One advantage of the graph representation of expressions is space efficiency. But only when this technique is combined with other optimisation techniques, performance issues will become evident. When an expression is translated to machine code, each subexpression is typically assigned one (pseudo) register of the underlying

Figure 7.3: An Expression $a \cdot a + a \cdot a$ as a Tree and as a Graph

machine. The fewer subexpressions there are, the less registers will be used and the less computations will be performed, and the easier it will be to schedule the use of the registers.

Cached Evaluation

During the transition instance analysis process described in Section 7.1, the same expressions will be evaluated several times, and typically only one variable will have changed its value between two successive evaluations of an expression.

One way to speed up the expression evaluator is to store the result of the last evaluation in a cache and return the cached value when the evaluator is called again. There is only one problem. When the value of a variable is changed in the assignment, the values of all expressions and subexpressions that depend on the value of the variable must be wiped out of the cache.

The overhead involved with managing the cache might over-weigh the advantage gained from using the cache. In [41], the cache can be disabled as a compile-time option. With the systems analysed so far, the expression evaluator has performed better with the cache enabled. Part of this can be attributed to common subexpression elimination.

Other Optimisations

It would be tempting to devise transformations for optimising the expression evaluator further. In a student project of a compiler construction course, the author implemented, among others, a transformation that was able to optimise an expression like $a + c + b + (b + a) \cdot 3$ to $(a + b) \cdot 4 + c$. Implementing and especially testing this kind of optimisations consumes much time, and the gain is often marginal, since the structures affected by such optimisations may seldom occur in practice.²

7.2.3 Interpreting vs. Compiling

A computer program that implements an abstract machine that executes computations expressed in some input language is called an *interpreter*. The interpreter itself may be executed by another interpreter, but at the lowest level there are some physical circuits that actually perform the computations.

²This also means that thorough testing of such optimisation algorithms is essential: errors in them are not likely to be discovered with simple test cases.

Simulating an abstract machines involves an overhead. Each expression or instruction expressed in the input language must be translated to a sequence of computations on the underlying machine. If the same computations are repeated several times, maybe varying some parameters, the computations could be speeded up by translating the input language to a lower-level language that can be executed on a simpler and faster abstract machine, or directly on a physical machine.

Expanding the Multi-Set Sum

The multi-set sum operation, described in Section 6.4.3, provides a compact way of representing initial markings and modelling complex actions, such as broadcasting a message to a number of neighbours.

One of the reasons why this powerful operation is rare in reachability analysers are implementation difficulties. When a multi-set sum occurs in an initialisation expression, it can be evaluated in a straightforward way, since there are no unknown variables. Multi-set sums on cause much more headache when they occur on arc expressions, since both the summation condition and the summand may depend on unbound variables.

Originally, we implemented the multi-set sum as a genuine operation, and the sums were expanded during the transition instance analysis. This complicated the bookkeeping, since the set of symbolic tokens associated with a transition varied during the instance analysis. Furthermore, a multi-set sum whose summation condition depends on other variables than summation indices cannot be fully expanded at all times. In addition to maintaining a “processed” flag for each symbolic token handled by NEXTTOKEN, we had to keep track on the multi-set sum expressions that were skipped due to an undefined summation condition.

To simplify and to speed up the instance analysis, we decided to expand the multi-set sums already in the parsing stage. Such a static transformation trades memory space for speed, since the algebraic term representing the symbolic expansion of the sum has to be kept in the memory all the time and not only when the sum actually needs to be evaluated.

When the static expansion takes place, the only variables whose values are available are the summation indices. The summation condition, which selects the index values for which the summand is evaluated, may depend on other variables. The condition can be translated to a scalar multiplication by zero or one.

Let there be the basic sorts $s, s', s'', b \in \mathcal{S}_\beta$ with the respective carriers $\mathcal{D}_s^{\mathcal{A}}, \mathcal{D}_{s'}^{\mathcal{A}}, \mathcal{D}_{s''}^{\mathcal{A}} = \mathbb{I}$ and $\mathcal{D}_b^{\mathcal{A}} = \mathbb{B} = \{\perp, \top\}$. Let the terms $T \in \mathbf{T}_b^{\mathcal{S}^\mu}(\mathcal{V})$ and $U \in \mathbf{T}_{\mu(s')}^{\mathcal{S}^\mu}(\mathcal{V})$ have \mathcal{V} such that $x \in \mathcal{V}_s$. The multi-set sum term

$$\mathbf{sum}_{x,T,U}$$

is equivalent to a term containing a chain of multi-set union operations defined in Section 6.4.5, enumerating through all the values $d \in \mathcal{D}_s^{\mathcal{A}}$ and combining the terms

$$\mathbf{mul}_s(U_x^d, \mathbf{select}_{b,s'',s'}(T_x^d, \mathbf{constant}_0(), \mathbf{constant}_1()))$$

where T_x^d and U_x^d denote transformations of the terms T and U such that each occurrence of the variable x has been replaced with the term $\mathbf{constant}_d()$.

Our implementation [41] folds constants while transforming the terms, and it omits the scalar multiplication or the whole term if the multiplicity is one or zero.

Compiling Expressions

The most obvious performance bottleneck in [41] is the expression evaluator, which makes heavy use of dynamically allocated objects. The performance improves vastly when the expressions are not interpreted but translated to something that can be compiled to machine code.

The programming languages C [28] and C++ [27] are the most obvious choices for an intermediate language. We selected C, since it is simpler than C++. Our implementation translates terms of basic sorts to functions that evaluate them. Variables are represented with pointers to values, and the null pointer denotes an undefined variable.

All basic terms that may need to be evaluated during the reachability analysis are translated to functions, along with the necessary data type definitions and auxiliary functions. Everything is compiled and linked to a library that is dynamically loaded in the reachability analyser. The interface with the reachability analyser consists of a number of arrays, depicted in Table 7.1.

In addition to the auxiliary functions that are part of the interface, our translator generates functions for computing the successor and predecessor, also for structured values, and the required type conversion functions for structured values.

At the first sight, it seems to be possible to translate each basic term in our input language to a C expression. There are some exceptions, such as the selection operation of Section 6.5, which cannot easily be translated to an expression when there are more than two alternatives. Structured values are troublesome. The C programming language [28] does not allow constructor expressions except as part of a variable declaration. It seems that this can be circumvented by using the comma operator to create a compound expression. Figure 7.4 illustrates these problems and shows how they can be solved by pushing assignments inside the expressions.

Structured values are handled as pointers. The generated code does not perform explicit dynamic memory allocation. Instead, the caller must allocate the space for the evaluation result. If the expression contains a structured subexpression whose evaluation result requires temporary storage, our translator declares a local variable for storing the intermediate result.

No code is generated for multi-set terms. The only basic terms whose operands are multi-set terms are the comparison and cardinality operations defined in Section 6.4.7 and 6.4.8. These terms are translated to calls to the run-time environment (the static core of the reachability analyser) via a call-back function. Since multi-set terms contain basic terms, the reachability analyser and the dynamically loaded code may invoke each other.

Nested invocations require the generated code be re-entrant, meaning that each invocation of the code has an own copy of all non-constant data it processes. Strictly taken, the variables, modelled with a static pointer table to data initialised by the run-time environment, are breaking the rule, and therefore, simultan-

```

typedef struct {
    int a;
    char b;
} struct_t;

void f (void) {
    struct_t s1;
    struct_t s2 = { 1, 'a' };           /* ok: declaration */
    s1 = { 2, 'b' };                   /* error */
    (s1.a = 2, s1.b = 'b', s1);       /* ok: compound expression */
    s2.a = (s1.a = 2, s1.b = 'b', s1).a; /* ok */
    (s1.a = 2, s1.b = 'b', s1).b = s2.b; /* error: invalid lvalue */
    s1.a = 2, s1.b = 'b', s1.b = s2.b; /* ok */
}

```

Figure 7.4: Treating Structured Values in C

ous calls to the expression evaluator must share the same valuation. This is not a problem in our analyser, which processes one transition instance at a time.

We currently translate all less-than comparison operations of structured values to calls to a generated function. When a value is compared against a constant, invoking the function clearly is less efficient than performing a short-circuited sequence of basic comparisons on the component values. The translation to a function call also requires the constants be initialised. Most of them can be initialised at compile time, but due to a restriction of the C programming language, the initialisation of some union-valued constants must be deferred until the dynamic library has been loaded and an initialisation routine has been invoked. A more efficient translation that does not require these constants will be implemented soon.

The `longjmp` operation of the C run-time library provides a convenient way for aborting the expression evaluation on an error, since the function can be invoked anywhere in an expression. In contrast, the `return` statement cannot be embedded in an expression.

The convenience of `longjmp` does not come without a price. Setting up the jump context is a costly operation especially on processors having a large bank of registers. Typical Algebraic System Net models use relatively simple expressions, and the `setjmp` operations required before each invocation of the expression evaluator introduce a considerable overhead. It seems that handling all errors in conditional return statements would be a major improvement.

7.3 Managing the Reachability Graph

The limited amount of system memory is a major bottleneck in exhaustive reachability analysis. Algorithms for reachability analysis and model checking need to keep track on the states that have been explored. In that way, they can detect cyclic behaviour and limit the investigation of successors to truly new states.

Table 7.1: The Interface to the Compiled Expression Evaluator

Symbol	Contents
unsigned s[]	sizeof(\mathcal{D}) for each type \mathcal{D}
void(*v[])(unsigned,void*)	$o_{\mathcal{D}}^{-1}$ functions
unsigned(*n[])(void*)	$o_{\mathcal{D}}$ functions
unsigned(*l[])(void*,void*,unsigned)	comparison functions ($<_{\mathcal{D}}, \leq_{\mathcal{D}}$)
unsigned(*c[])(void*)	constraint checks ($d \in \mathcal{D}$)
void*x[]	valuation (v)
void(*e[])(jmp_buf*,void*)	expressions ($e_v^{\mathcal{A}}(T)$)
unsigned(*el[])(void*)	assignment candidates ($\langle x, d \rangle$)
enum Error (*ec[])(void*)	compatibility checks ($T \approx_v \mu$)
unsigned(**cb)(jmp_buf*,void*,void*)	call-back for multi-set terms

There are some techniques that only manage the set of reachable states and utilise similarities between the states. One of them, Binary Decision Diagrams [4, Chapter 5], has been successfully applied in the verification of digital circuits. Techniques applied on the analysis of software systems include a state compaction method for product automata [16] and a method known as Graph Encoded Tuple Sets [21].

One problem with these techniques is that inserting a state may involve global changes, slowing down disk-based implementations. Another problem is that states have no identities: there is no way to retrieve a state from the structure by specifying an index number. It is difficult to use such a structure for anything else than determining whether a particular state has been explored.

Since we want to be able to make all sorts of queries on the generated reachability graph, our approach represents each marking (state) separately and records also the transition instances that lead from one marking to another. Our approach encodes the markings in a compact way, and it saves system memory by maintaining the encoded markings and transition instances on disk.

The reachability graph storage implemented in [41] builds heavily on two encoding routines. One inputs a number in a set $\{0, \dots, n-1\}$ and appends a string of $\lceil \log_2 n \rceil$ binary digits to an encoding buffer. Another routine encodes a value $d \in \mathcal{D}^{\mathcal{A}}$ using around $\lceil \log_2 |\mathcal{D}^{\mathcal{A}}| \rceil$ binary digits.³ This is similar to the approach represented in [13], but it was developed independently.

Our encoding is optimal if $n = 2^k$ for some integer $k \geq 0$ and if each item of \mathcal{D} occurs with equal probability. More efficient encodings would be possible if the data to be encoded were known in advance. For a parameterised model, one could try to estimate probabilities for different values by analysing smaller state spaces generated with suitable parameters. We did not consider this option further.

³Due to an implementation choice, if $|\mathcal{D}^{\mathcal{A}}|$ cannot be represented in a machine word, the value will be encoded componentwise, which wastes a fraction of a binary digit for each component whose domain is not of size 2^k for some k .

7.3.1 Encoding Markings

The marking M of a net $\mathcal{N} = \langle \mathcal{P}, \mathcal{T}, \mathcal{F} \rangle$ is encoded by processing the places in a systematic order. For each place $p \in \mathcal{P}$, the local marking

$$\mu = M(p)$$

must be encoded. With the assumption that $\mu : \mathcal{D} \rightarrow \mathbb{N}$ maps most data items in \mathcal{D} to zero, it makes sense to explicitly represent the data items with nonzero multiplicity. Obviously, each data item can be represented using $\lceil \log_2 |\mathcal{D}| \rceil$ binary digits. Therefore, we shall concentrate on the encoding of the multiplicities and the representation of empty places.

Representing Multiplicities

A multi-set μ can be characterised by two quantities: the total number of items

$$t = |\mu|$$

and the number of distinct items

$$d = |\{a \mid \mu(a) > 0\}|.$$

The total number of items can theoretically be any natural number, but a finite-memory implementation imposes a limit on it, typically $0 \leq t < 2^n$ for some n .

An user-defined *capacity constraint* [41] can reduce the number of bits required for representing t . If there are m different possibilities for the total number of tokens in a place, the actual number t can be represented using $\lceil \log_2 m \rceil$ bits.

Encoding the total number of items t before the number of distinct items d has one advantage: it is straightforward to see that $1 \leq d \leq t$ when t is nonzero. Therefore, d can be represented using $\lceil \log_2 t \rceil$ bits. After this, the distinct items

$$\{a \mid \mu(a) > 0\}$$

and their multiplicities $\mu(a)$ are encoded in descending order of $\mu(a)$. Clearly

$$\left\lceil \frac{t}{d} \right\rceil \leq \mu(a) \leq 1 + t - d$$

holds for the greatest multiplicity $\mu(a)$. If $\mu(a) = 1 + t - d$, it holds that the other $d - 1$ distinct items must have a multiplicity of 1 in order for the total number of items to be t . Similarly, if $\mu(a) = \lceil \frac{t}{d} \rceil$, the multiplicities of the remaining items must equal $\mu(a)$ or $\mu(a) - 1$.

So, the greatest multiplicity $\mu(a)$ can always be represented with

$$\left\lceil \log_2 \left(2 + t - d - \left\lceil \frac{t}{d} \right\rceil \right) \right\rceil$$

binary digits. After decoding $\mu(a)$, the decoder knows the remaining total cardinality $t' = t - \mu(a)$ and the number of remaining distinct items $d' = d - 1$. When

the multiplicities are encoded in descending order, the encoder always selects the greatest of the remaining multiplicities and uses less and less bits.

This encoding of multiplicities appears to be quite compact even when capacity constraints are not used. A simpler encoding might represent the multiplicities using a fixed number of binary digits for the multiplicity of each distinct item. For instance, when the multiplicities $\mu(a)$ are limited in the range $0 \leq \mu(a) < 2^n$ for some n , such a simple coding would use up at least dn bits for encoding d multiplicities, not considering the bits needed for signalling the end of the encoded stream.

For representing $d = 5$ multiplicities, the simple encoding would use at least $5n$ bits. The optimised encoding needs n bits for representing the total cardinality. Assuming that it is 8, the number of distinct tokens is encoded in 3 bits. The greatest multiplicity lies between $\lceil \frac{8}{5} \rceil = 2$ and $8 - 5 + 1 = 4$; therefore it can be represented with 2 bits. Clearly, the improved encoding requires less than $n + 3 + 5 \cdot 2 = n + 13$ bits. The difference between $5n$ and $n + 13$ is tangible in practical implementations, which typically use $n = 16$ or $n = 32$.

Representing Empty Places

In many practical models, there is a substantial number of empty places in most reachable markings. With our optimised multiplicity encoding, an empty place requires $\lceil \log_2 m \rceil$ bits of storage, if there are m different possibilities for the total number of tokens in the place.

As it is rather uncommon to define tight capacity constraints, representing the total cardinalities typically requires one machine word per place, which easily dominates the encoding of markings where most places are empty. There ought to be a more compact encoding for the empty places.

Our solution is to encode the number e of empty places in the marking, $0 \leq e \leq |\mathcal{P}|$. This requires $\lceil \log_2(|\mathcal{P}| + 1) \rceil$ binary digits. This number alone does not provide enough information for the decoder, unless it happens to be 0 or $|\mathcal{P}|$. In the following, we assume that $0 < e < |\mathcal{P}|$.

If $e > \frac{1}{2}|\mathcal{P}|$ —that is, there are more empty places than nonempty ones—then it makes sense to represent the identity of the nonempty places. Our implementation identifies each place $p_i \in \mathcal{P}$ with an index number $0 \leq i < |\mathcal{P}|$, and it encodes the markings of the places in ascending order of index numbers.

Clearly, the smallest index number i_1 of a nonempty place must be in the range

$$0 = l_1 \leq i_1 \leq h_1 = e$$

since there are at most e empty places in the beginning of the sequence. So, i_1 , the index number of the first nonempty place, can be stored using $\lceil \log_2(h_1 - l_1 + 1) \rceil$ binary digits. What about the following nonempty places i_{k+1} ? It holds that

$$i_{k+1} \geq l_{k+1} = i_k + 1,$$

since the indices are encoded in ascending order. It is easy to see that there are $i_k - (k - 1)$ empty places before i_k , since i_k is the k th smallest index of a nonempty

place. Thus, of the remaining places, $e - (i_k - (k - 1))$ are empty, and for the upper limit $h_{k+1} \geq i_{k+1}$ we have

$$\begin{aligned} h_{k+1} &= l_{k+1} + e - (i_k - (k - 1)) \\ &= i_k + 1 + e - i_k + k - 1 \\ &= e + k. \end{aligned}$$

Since $h_1 = e$, it is easy to see that $h_{k+1} = h_k + 1$.

Similarly, if the empty places are a minority, we represent only the indices of empty places. This is completely analogous to the previous case; we just let $h_1 = |\mathcal{P}| - e$ in the beginning.⁴

Implicit Places

Some models contain *implicit places*, places whose local marking is a function of the local markings of other places. One might ask why such places occur in models, but they can make models more readable and ease the transition instance analysis. Nevertheless, the local markings of these places need not be encoded. Allowing the user to specify marking-dependent initialisation expressions would be an elegant way to point out implicit places.

This optimisation is likely to be implemented soon in [41]. It, together with the capacity constraint, slightly affects the formal semantics: when the firing of an enabled transition instance would violate a capacity constraint or be in contradiction with the initialisation expression of an “implicit place,” the transition instance would not be fired, but it would be reported to be faulty.

Hashing

A hash value will be computed for the encoded marking. If the hash value does not exist in a *hash table*, a memory-based data structure that maps hash values to numbers of corresponding markings, the encoded marking will represent a new node in the reachability graph. In this case it will be assigned its own number and stored in the hash table as well as written to a disk file containing the distinct encoded markings generated so far.

If the hash value exists in the hash table, the encoded markings corresponding to it will be fetched from the disk file and compared against the newly encoded marking. If a match is found, the newly encoded marking will be assigned the number of the existing marking. Otherwise the marking will be written to the disk file and added to the hash table.

7.3.2 Encoding Transition Instances

The encoded reachability graph includes also transition instances, leading from one marking to another. Each transition instance consists of a high-level transition name and an assignment, mapping variables to values.

⁴Of course this makes a big difference in our implementation, where the indices are interleaved with the markings in the encoded stream.

Each high-level transition $t \in \mathcal{T}$ is assigned a unique number between 0 and $|\mathcal{T}| - 1$. The encoded instance consists of a pair of marking numbers (the source and the target marking), followed by the transition number encoded in $\lceil \log_2 |\mathcal{T}| \rceil$ binary digits and by the encoded assignment.

The assignment is encoded by processing the variables in a systematic order. One bit is used for signalling that a variable is undefined, if it is allowed to be undefined.⁵ When a variable is defined, its encoded value will be appended to the bit vector representing the assignment.

No search structure for transition instances is needed while generating the reachability graph. The transition instances can be appended to a disk file in a linear manner. An application that needs to refer to the transition instances can save memory and time by reading the disk file in one pass and constructing a data structure that contains the numbers of the source and target marking and an offset to the encoded transition instance in the disk file.

Some resources could be saved by allowing the user to specify the variables whose values should be included in the encoded transition instance. The omitted values could be reconstructed on demand by running the instance analysis algorithm in the source marking of the event.

7.4 MARIA as a Software Project

The beginning of this chapter describes our reachability analyser implementation [41] from a purely theoretical perspective, not mentioning the low-level details. Designing and implementing the reachability analyser has taken the author over one year of intensive work. Developing a piece of software that consists of more than one million characters or 48,000 lines of source code—excluding the automatically generated parts—is not an easy task. We shall now describe some of the methods and tools that helped us to ensure high quality at all stages of the project.

7.4.1 Development Tools

Choosing the Programming Language

One of the major design constraints of the reachability analyser is portability. It must be possible to run the source code through any compliant compiler, and the resulting executable must work properly, no matter what the underlying computer architecture looks like.

There are not very many truly portable non-interpreted languages to choose from. Due to the size of the project, special attention had to be paid to the maintainability of the code. Since the analyser is being extended with new algorithms by a number of people, it is important to keep the source code as modular as possible. It is possible to write modular programs in traditional procedural languages

⁵By default, all transition variables must have a value in all enabled transition instances. Undefined variables occur when the multiplicity of an input arc expression evaluates to zero.

such as C [28], but the object-oriented approach has language constructs for coupling data structures and code together and for specifying the hierarchy and the internal interfaces of a program.

C++ is probably the most widely spread object-oriented language that can be compiled to machine code. Since its roots are in the C programming language, it suits well also to very low-level programming. C++ has some drawbacks: Since the language is extremely complex, not all presently available C++ compilers support all the features completely, and there are many ways to “shoot oneself in the foot.” Many style guides list a number of C++ features that one should refrain from using in order to avoid all kinds of unexpected problems.

For us, the advantages of C++ weighed more in the scale than the disadvantages: Since C++ is available for a wide variety of computer platforms, it has good portability. Also the Standard Template Library, a collection of well-tested algorithms for managing basic data structures, has probably saved us days of programming and debugging. It lets the programmer view the data structures as containers, “black boxes” with clean interfaces for storing, retrieving and removing elements.

The C++ Compiler

Since the C++ programming language has evolved a lot during its history and has only recently established itself as an international standard [27], there are not so many compilers that support all features of the language. We had access only to two compilers that satisfied our needs: the GNU Compiler Collection [54] and the Digital C++ Compiler [8].

The Hewlett–Packard compiler [23] we had access to managed to compile some of the code, but we could not circumvent all compiler or linker bugs. At some stage, the automatically generated parser functions appeared to be too big for the compiler. By replacing the break statements in the parser switch block with equivalent goto statements, we could process the parser module through the compiler. At that time, it was possible to compile and link a heavily stripped or modified version of the program, but the executable code did not work properly. At the present stage of development, trying to use the Hewlett–Packard compiler appears to be a hopeless task, since the compiler silently fails on at least one input file, exiting with a nonzero exit code without producing any diagnostic output.

Luckily, the GNU Compiler Collection is available for a wide variety of computing platforms, and therefore also our reachability analyser can be used on several hardware platforms. Our main development platform is Linux [58]; we mainly use the Digital compiler for ensuring portability, and for detecting some thinking errors, such as making implicit assumptions about the machine word length.

The reachability analyser makes use of some C++ features that were not part of the original language [56], such as the Standard Template Library, an algorithm library for basic data structure management. For input and output, we use the `<stdio.h>` library of the C programming language [28], since we feel that using the `<iostream>` library advertised in practically every C++ textbook would make the code more complicated and inefficient. For similar reasons, we decided not to use

run-time type identification or exception handling.

Parser Construction Tools

Since constructing a parser for a nontrivial language is a challenging task, we did not even consider writing parsers by hand. When one has the liberty of developing the language and the parser simultaneously, it is easy to avoid constructs that are difficult to parse. An advanced parser generator such as Bison [6] detects all potential ambiguities in the grammar, some which can be quite complicated. Handling errors in the input can substantially complicate a hand-written parser. Bison also provides a relatively easy and safe way for handling erroneous input.

The reachability analyser inputs the model to be analysed and its desired properties in a language whose grammar is defined with a Backus–Naur Form (BNF) grammar. For this class of languages, the *de facto* parser construction tools are Flex [47], a lexical analyser generator based on deterministic finite automata, and Bison, an LALR(1)⁶ parser generator. Although these tools have been designed to generate C programs, the code they produce can be compiled with a C++ compiler, and it is possible to write the user-supplied code, the so called *semantic actions* in C++. This is exactly how the parsers of our reachability analyser were developed. We decided not to use the Flex option that would generate a C++ class of the scanner code, since Bison lacks a similar option.

Version Control

Often there are cases where one makes a modification to a file and regrets it later, wanting to revert to an earlier version of the file. The problem could be solved by backing up each file before it is modified, but this solution is inefficient if the modifications are small compared to the size of the file. A *version control system* is a much better choice.

A version control system records the modification history of files. When a file is kept in a version control system, it is possible to retrieve any previously stored version of the file and to see the differences between any versions.

Sometimes a software developer may want to try an alternative algorithm to solve a problem. Usually he just wants to disable the existing algorithm e.g. by enclosing it in comment delimiters or by using a preprocessor directive. If the old algorithm turns out to be better than the new one, the new algorithm can be deleted and the old code enabled again. When version control is used, there is no need to leave large sections of disabled code in the program, and the program code remains compact and easily understandable. If an error is encountered in the new algorithm, the old algorithm can be restored from the version control system.

There are two freely available version control systems: RCS [57] and CVS [2]. We chose CVS, because it has been designed for distributed setups and does not require files to be locked for modification. The database, called *repository*, is kept in one place, and files can be retrieved from it even across the network. Even when

⁶LALR(1) is the class of languages that can be parsed from Left to Right with a Look-Ahead of 1 symbol. [1]

the development takes place on several workstations, only the repository needs to be backed up. In the case of a disk crash, only those modifications that had not yet been checked in to the repository will be lost.

Integrated Development Environment

GNU Emacs [12] is one of the most versatile interactive text processing programs available. The editor has a built-in interpreted programming language, in which most of its functions have been programmed. For various languages, there are editing modes, which know something about the lexical rules or the grammar of the language. Built-in indentation rules make it easy to format code in an aesthetically pleasing way, and syntax highlighting—that is, multi-colour display of keywords and various language constructs—causes addiction easily.

To make our reachability analyser friendlier to the user, we developed a GNU Emacs mode for our modelling language. Keywords of the language as well as names of data types, places and transitions are highlighted appropriately, and the common indentation rules for the C programming language work pretty well with our modelling language.

7.4.2 Selected Implementation Details

Many things that appear to be simple or trivial from a purely theoretical point of view may actually require most of the attention when constructing a software implementation.

The Modelling Language

One of the implementation details that we have deliberately excluded from our theoretical discussion is the syntax, the actual appearance of the modelling language understood by the analyser. A detailed grammar description is available in [41]; here we shall briefly discuss some of the design choices that were made.

There were two main ideas in the design of the language: there shall be no forward references, and apart from that limitation, the language should be as flexible to use as possible. The parser is relatively easy to implement in the absence of forward references, since everything can be transformed into the final syntax tree in one pass.

If a language allows definitions to be used or referenced before their definition, the parser must store at least some language constructs in an intermediate form. Clearly, all unresolved references have to be literally stored in the intermediate syntax tree. But also the line numbers and source file names of such references have to be recorded, so that diagnostic messages can be displayed for unresolved or contradictory references once the whole input has been parsed. When there are no forward references, most diagnostic messages can be issued “on the spot,” and the file names and line numbers can largely be omitted from the syntax tree.

A language with no forward references can still be rather flexible. There is no need to further limit the order definitions may occur in the input. Although one

may consider it good practice to define all data types in the beginning of a model description, followed by the place definitions and finally the transition definitions, it would be too restrictive to enforce such a practice in the parser. Our implementation even allows transitions to be defined in several fragments, a few arc expressions at a time.

Name spaces are part of the flexibility. Our modelling language has been designed so that names of places, transitions, variables of each transition, data types, and enumeration constants of each enumerated data type, all have their own name spaces. Although one may consider it bad practice to use the same name for several entities of different type, it is possible to do so. Unintentional mistakes can be pointed out by specifying extended regular expression [29] patterns for the names allowed or disallowed for different entities.

The preferred low-level syntax of a language is a matter of personal taste. Our expression syntax resembles the C notation, which—in our opinion—should make it easier to learn the modelling language. Some tool authors prefer a mathematical-looking notation and may use e.g. pairs of the slash and backslash symbols to represent the logical conjunction and disjunction operators. We prefer to use the backslash as an escape character, in accordance with a common practice expressed e.g. in [29].

The Data Type System

Chapter 5 describes the data type system implemented in [41] from the user point of view. Internally, each data type construct is implemented as two abstract C++ classes, one for the data type itself and one for its values. There are a few common methods for all data types, and some specialised methods for each particular data type. The common methods are mainly related with type compatibility checking, type conversions and the ordering of values.

Data type compatibility checking is implemented as two methods that compute an asymmetric relation between types. The stronger variant of type compatibility requires that all values of the first type are values of the second type, while the weaker variant only requires that there exist compatible values in the first type. When only the weak compatibility holds, a third method is used for determining whether a specific value of the first type is compatible with the constraint of the second data type.

The data type conversions described in Section 6.3.6 are implemented in one rather straightforward virtual method. Implementing the bijective mapping $o_{\mathcal{D}} : \mathcal{D} \rightarrow \{0, \dots, |\mathcal{D}| - 1\}$ defined in Section 5.1.1 deserves more attention. While it is theoretically possible to implement the mapping by applying a series of operations computing successors or predecessors, there exists a much faster solution. The system keeps track on the number of possible values $|\mathcal{D}|$ for each data type \mathcal{D} , also for constrained data types. The structured data types defined in Section 5.3 are handled componentwise, utilising equalities implied by the lexicographical ordering.

There is one restriction in the implementation of $o_{\mathcal{D}}$: structured data types with $|\mathcal{D}|$ greater than the capacity of a machine word cannot be converted to numbers

and vice versa. Therefore, the reachability graph encoder described in Section 7.3 manages these types as special cases, performing the encoding componentwise. The restriction actually bites the user at two points: such huge data types cannot be used for indexing arrays⁷ nor as the left-hand side of a selection operator for short-circuit evaluation.

Other methods related with the ordering of values implement comparisons and the successor and predecessor relations. They have been implemented so that all operations agree on the total order defined by the methods that compute $o_{\mathcal{D}}$.

⁷This is a very marginal restriction, since the smallest meaningful array—a bit array—with 2^n indices would require 2^n bits of storage, more than the virtual memory capacity of most computers.

Chapter 8

Constructing and Analysing Models

A formalism without any practical applications is a dead formalism. We will present two examples that illustrate some of the advantages of the algebraic operations presented in Chapter 6 over the set of operations conventionally defined in computer tools for high-level Petri Nets.

8.1 Point-to-Multipoint Communications

Many distributed algorithms designed for computer networks involve operations where a node sends messages to or waits for a message from each or some of its neighbours. When constructing parameterised models, e.g. with n nodes, the multi-set sum operator defined in Section 6.4.3 provides a compact way for writing arc expressions describing such situations.

Consider a system of n nodes connected via a broadcasting network. The nodes periodically need to synchronise with each other. This is accomplished by broadcasting messages. When a node is ready for synchronisation, it will send a message to all other nodes and wait for a message from all other nodes. In order to model this system, we need to define the neighbourhood relation

$$N = (U \times U) \setminus \{\langle d, d \rangle \mid d \in U\}$$

with $|U| = n$.

Without using the multi-set sum operator, this relation, or its projection

$$N(u) = \{u' \mid \langle u, u' \rangle \in N\}$$

can only be represented for some fixed value of n . Let us consider an example where $n = 4 = |U|$. Let \mathcal{A} be an algebra with $\mathcal{D}_s^{\mathcal{A}} = U$ and $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{T}_{U,U}$. Now

$$\begin{aligned} N(u) &= \left\{ u' \mid \langle u', 1 \rangle \in e_{\{(x,u)\}}^{\mathcal{A}}(T_1) \right\} \\ T_1 &= \mathbf{union}_s(\mathbf{mset}_s(\mathbf{succ}(x), \mathbf{constant}_1()), \\ &\quad \mathbf{mset}_s(\mathbf{succ}(\mathbf{succ}(x)), \mathbf{constant}_1()), \\ &\quad \mathbf{mset}_s(\mathbf{succ}(\mathbf{succ}(\mathbf{succ}(x))), \mathbf{constant}_1())). \end{aligned}$$

It is easy to see that in this kind of a construct, the length of the term T_1 presented above increases at least linearly with n . The multi-set sum operator allows for more compact notation, which does not depend on n , the number of items in the set U :

$$\begin{aligned} T_1 &= \mathbf{sum}_{y, T_2, T_3} \\ T_2 &= \mathbf{unequal}(x, y) \\ T_3 &= \mathbf{mset}_s(y, \mathbf{constant}_1()). \end{aligned}$$

Now we are ready to represent an Algebraic System Net model of the system, illustrated in Figures 8.1 and 8.2. The latter uses graphical notation, representing places (elements of P) with circles encircling sort identifiers, transitions (elements of T) with boxes, and the flow relation F with directed arcs connecting the images of places and transitions. Furthermore, the inscription function i is represented with textual inscriptions written next to the places, transitions and arcs. Trivial “null” inscriptions—empty place initialisers and constantly enabled transition guards—are omitted. We have also labelled the place and transition symbols with the names used in the purely textual representation of Figure 8.1.

The alert reader may notice that the place **PENDING** in Figure 8.1 is redundant. In all reachable markings of the net, the place **BUS** contains a set of tuples

$$\bigcup_{x \in A} \{\langle x, y \rangle \mid y \in B \setminus \{x\}\}$$

where A and B are some subsets of $\mathcal{D}_s^{\mathcal{A}}$. Due to the structure of the net, the place **PENDING** will contain the set A corresponding to the contents of the place **BUS** in all reachable markings.

The redundant place **PENDING** is needed in the model to circumvent the limitations of the transition instance analysis algorithm described in Section 7.1.1. In our algebraic terms, if we allow initialisation expressions of implicit places to depend on the current marking of the net as suggested in Section 7.3.1, the local marking of the redundant place can be represented as follows:

$$\begin{aligned} i(\mathbf{PENDING}) &= \mathbf{mmap}_{x, y, T}(\mathbf{map}_{z, U}(\mathbf{BUS})) \\ T &= \mathbf{constant}_1() \\ U &= \mathbf{component}_{s', 1}(z). \end{aligned}$$

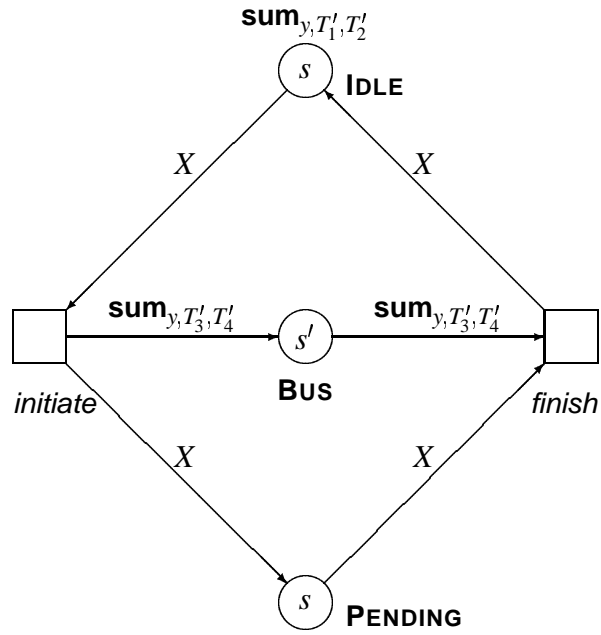
It is also easy to see that the place **IDLE** is the complement place of **PENDING**. In all reachable states, the union of the two local markings equals $\mathcal{D}_s^{\mathcal{A}}$. In other words,

$$\begin{aligned} i(\mathbf{IDLE}) &= \mathbf{minus}_s(\mathbf{sum}_{x', T', U'}(\mathbf{PENDING})) \\ T' &= \mathbf{constant}_{\top}() \\ U' &= \mathbf{mset}_s(x', \mathbf{constant}_1()). \end{aligned}$$

It is not always clear whether a place is implicit. The theory for finding implicit or redundant places is called *invariant analysis*, but it is not with the scope of this work. However, it may be noted that our tool [41] is capable of determining whether certain place invariants, as defined in [37], hold in the reachable states of a model.

$$\begin{aligned}
\Sigma &= \langle \mathcal{N}, \mathcal{A}, \mathcal{V}, i \rangle \\
\mathcal{N} &= \langle \mathcal{P}, \mathcal{T}, \mathcal{F} \rangle \\
s, s' &\in S_{\beta} \\
\mathcal{D}_s^{\mathcal{A}} &= \{0, \dots, n-1\} \\
\mathcal{D}_{s'}^{\mathcal{A}} &= \mathbb{T}_{\mathcal{D}_s^{\mathcal{A}}, \mathcal{D}_s^{\mathcal{A}}} \\
\mathcal{P}_{\mu(s)} &= \{\mathbf{IDLE}, \mathbf{PENDING}\} \\
\mathcal{P}_{\mu(s')} &= \{\mathbf{BUS}\} \\
\mathcal{P} &= \mathcal{P}_{\mu(s)} \cup \mathcal{P}_{\mu(s')} \\
\mathcal{T} &= \{\mathit{initiate}, \mathit{finish}\} \\
\mathcal{F} &= \{f_1, f_2, f_3, f_4, f_5, f_6\} \\
f_1 &= \langle \mathbf{IDLE}, \mathit{initiate} \rangle \\
f_2 &= \langle \mathit{initiate}, \mathbf{PENDING} \rangle \\
f_3 &= \langle \mathit{initiate}, \mathbf{BUS} \rangle \\
f_4 &= \langle \mathbf{PENDING}, \mathit{finish} \rangle \\
f_5 &= \langle \mathbf{BUS}, \mathit{finish} \rangle \\
f_6 &= \langle \mathit{finish}, \mathbf{IDLE} \rangle \\
\mathcal{V} &= \mathcal{V}_s \\
\mathcal{V}_s &= \{x\} \\
i(\mathbf{IDLE}) &= \mathbf{sum}_{y, T'_1, T'_2} \\
i(\mathbf{PENDING}) &= \mathbf{empty}_s() \\
i(\mathbf{BUS}) &= \mathbf{empty}_{s'}() \\
i(\mathit{initiate}) &= \mathbf{constant}_{\top}() \\
i(\mathit{finish}) &= \mathbf{constant}_{\top}() \\
i(f_1) &= \mathbf{mset}_s(x, \mathbf{constant}_1()) \\
i(f_2) &= \mathbf{mset}_s(x, \mathbf{constant}_1()) \\
i(f_3) &= \mathbf{sum}_{y, T'_3, T'_4} \\
i(f_4) &= \mathbf{mset}_s(x, \mathbf{constant}_1()) \\
i(f_5) &= \mathbf{sum}_{y, T'_3, T'_4} \\
i(f_6) &= \mathbf{mset}_s(x, \mathbf{constant}_1()) \\
T'_1 &= \mathbf{constant}_{\top}() \\
T'_2 &= \mathbf{mset}_s(y, \mathbf{constant}_1()) \\
T'_3 &= \mathbf{unequal}(x, y) \\
T'_4 &= \mathbf{mset}_{s'}(\mathbf{cons}_{s'}(x, y), \mathbf{constant}_1()).
\end{aligned}$$

Figure 8.1: Algebraic System Net Model of a Synchronisation Protocol



$$\begin{aligned}
 s, s' &\in \mathcal{S}_\beta \\
 \mathcal{D}_s^{\mathcal{A}} &= \{0, \dots, n-1\} \\
 \mathcal{D}_{s'}^{\mathcal{A}} &= \mathbb{T}_{\mathcal{D}_s^{\mathcal{A}}, \mathcal{D}_s^{\mathcal{A}}} \\
 \mathcal{V} &= \mathcal{V}_s \\
 \mathcal{V}_s &= \{x\} \\
 X &= \mathbf{mset}_s(x, \mathbf{constant}_1()) \\
 T'_1 &= \mathbf{constant}_\top() \\
 T'_2 &= \mathbf{mset}_s(y, \mathbf{constant}_1()) \\
 T'_3 &= \mathbf{unequal}(x, y) \\
 T'_4 &= \mathbf{mset}_{s'}(\mathbf{cons}_{s'}(x, y), \mathbf{constant}_1())
 \end{aligned}$$

Figure 8.2: Graphical Representation of the Model in Figure 8.1

8.2 Existential Quantification

In some models one would like to apply existential and universal quantification on algebraic terms. For example, minimisation problems, that is, problems of the form

$$\exists x : \forall y : f(x) \leq g(x, y),$$

need to be solved every now and then. One solution to this kind of problems is the introduction of an algebraic operation; another way to attack the problem is to introduce implicit places, e.g. a place that holds the current minimum.

Purely existential problems can be solved by counting. If we want to know whether there exists a token i in a place p such that some logical property $f(i)$ holds, we can introduce an implicit place that holds the number of such tokens i in the place p , in all reachable markings. A similar approach can be taken for problems involving purely universal quantification, since $\exists i : f(i)$ is equivalent to $\forall i : \neg f(i)$.

An existential-universal quantification operation has been considered for implementation in [41], but it has not been implemented yet. One problem with such quantification is that if both quantification variables x and y belong to a data type with n items, the algorithm will need to test n^2 combinations in the worst case. On the other hand, performing the quantification in one atomic step is an order of magnitude more efficient than modelling it “manually” with a sequence of transitions, causing numerous uninteresting intermediate states to be generated.

8.3 The Performance of Exhaustive Analysis

We shall demonstrate the performance of our reachability analyser implementation with a small example. The distributed data base system model presented in Figure 8.3 has been translated into the input language of our tool from an example file distributed with PROD [62].

The model has one parameter, the number of service nodes. The original model for PROD, dbm.net, describes a system with ten nodes. Adding or removing nodes involves changes in several places of the model. The MARIA model refers to the number of service nodes only in the definition of the data type `db_t`. The multi-set summations used in the initialisation expression of the place **INACTIVE** and in some arc expressions expand according to the domain size of this data type.

8.3.1 The Size of the Encoded State Space

Table 8.1 illustrates the performance of our state encoding on a 32-bit computer system. We analysed the distributed data base model with PROD using the default options, and with MARIA using two variants of the model: without and with capacity constraints for the places.

The figures include the space required for bookkeeping. In MARIA, the bookkeeping record is a table of hash values and file offsets. On a 32-bit system with

```

typedef unsigned (1..10) db_t;
typedef struct {
    db_t sender;
    db_t recipient;
} db_pair_t;

place INACTIVE (0..#db_t) db_t: db_t d: d;
place WAITING (0..1) db_t;
place PERFORMING (0..#db_t) db_t;
place EXCLUSION (0..1) struct {}: {};
place SENT (0..#db_t) db_pair_t;
place RECEIVED (0..#db_t) db_pair_t;
place ACKNOWLEDGED (0..#db_t) db_pair_t;

trans "update and send messages"
in { INACTIVE: s; EXCLUSION: {}; }
out { WAITING: s; SENT: db_t r (r != s): { s, r }; };

trans "receive acknowledgements"
in { WAITING: s; ACKNOWLEDGED: db_t r (r != s): { s, r }; }
out { INACTIVE: s; EXCLUSION: {}; };

trans "receive message"
in { INACTIVE: r; SENT: { s, r }; }
out { PERFORMING: r; RECEIVED: { s, r }; }
gate s != r;

trans "send acknowledgement"
in { PERFORMING: r; RECEIVED: { s, r }; }
out { INACTIVE: r; ACKNOWLEDGED: { s, r }; }
gate s != r;

```

Figure 8.3: MARIA Model of a Distributed Data Base System

Table 8.1: Encoded State Space Sizes for the Distributed Data Base Model

Model Size		Encoded State Space in Bytes			Bytes per State		
Nodes	States	PROD	MARIA		PROD	MARIA	
1	2	167	36	24	83.5	18	12
2	7	407	164	84	58.14	23.43	12
3	28	1,625	788	360	58.04	28.14	12.86
4	109	7,497	3,524	1,740	68.78	32.33	15.96
5	406	34,595	14,920	6,492	85.21	36.75	15.99
6	1,459	154,739	55,292	29,176	106.06	37.90	20.00
7	5,104	668,709	217,360	102,076	131.02	42.57	20.00
8	17,497	2,799,605	753,432	415,792	160.00	43.06	23.76
9	59,050	11,396,735	2,797,476	1,457,512	193.00	47.37	24.68
10	196,831	45,271,215	9,365,904	5,511,256	230.00	47.58	28.00

Table 8.2: State Space Generation Times for the Distributed Data Base Model

Model Size		Original		Optimised		Compiled	
Nodes	States	user/s	sys/s	user/s	sys/s	user/s	sys/s
1	2	0.0	0.0	0.0	0.1	0.7	0.1
2	7	0.0	0.0	0.0	0.0	0.9	0.1
3	28	0.0	0.0	0.0	0.0	1.1	0.1
4	109	0.1	0.0	0.1	0.0	1.4	0.1
5	406	0.6	0.1	0.4	0.1	1.6	0.1
6	1,459	3.1	0.2	2.0	0.3	2.8	0.2
7	5,104	13.3	1.0	8.8	0.9	6.9	1.1
8	17,497	56.1	4.3	36.7	4.4	23.9	4.2
9	59,050	229.3	19.3	149.1	18.1	89.5	18.3
10	196,831	902.3	73.6	591.5	71.3	349.1	68.3

32-bit file offsets, the bookkeeping overhead is 8 bytes per encoded state. The low-level binary string routines in MARIA operate on machine words, and the encoded states are stored as sequences of machine words. Thus, on a typical system, the encoded state occupies at least three machine words—12 bytes in Table 8.1.

Our encoding algorithm outlined in Section 7.3.1 utilises the capacity constraint when encoding the total number of tokens in a place. When there is no capacity constraint, our implementation represents the number using one machine word. The other extreme is a capacity constraint that allows only one value for the total number of tokens. The first line of Table 8.1 confirms that defining a capacity constraint can save up to one machine word for each non-empty place.

8.3.2 Execution Times of Exhaustive Reachability Analysis

Users are not only after advanced features and algebraic operations; also practical things, such as performance, are important for them. Table 8.1 presents the execution times the analyser requires for exhaustively analysing the distributed data base model. There are two subcolumns: time spent in user mode, and time required for operating system services, such as low-level file operations and memory allocation. These figures were obtained on an Intel Pentium II system running at 266 MHz, using the GNU Compiler Collection [54] 2.95.2 and Linux [58] 2.2.7.

Improvements

Our original implementation expanded multi-set sums during transition instance analysis. This complicated the algorithms considerably, and it prohibited a translation of the expressions into machine code, since the expanded expressions did not exist prior to transition instance analysis. The second and the third column of Table 8.1 show how implementing the expression transformations described in Section 7.2.3 affected the performance.

It is a little surprising to see how little invoking the C compiler costs. Of course, the bigger the model, the bigger the generated C program becomes. On the other hand, the expressions in this simple model are very short, and there is not much difference between interpreted evaluation and direct execution of machine code.

Probably most of the time savings in the analysis of this model can be attributed to memory management. The run-time environment for the generated code does not allocate memory from the heap except when inserting items to multi-sets.

A New Bottleneck

Currently the biggest bottleneck in the reachability analyser is the operation that inserts an item d to a multi-set μ . The processor spends about a tenth of the total execution time performing these operations.

The multi-set is implemented as a search structure of $\langle d, \mu(d) \rangle$ pairs with $\mu(d) > 0$. The insertion operation traverses the search structure, comparing items using a compiled function. The less-than comparison it applies is rather expensive, especially for complex structured values.

We are planning to generate hash functions for each data type. Then the less-than comparisons could be replaced with much cheaper equality comparisons, which would only be invoked when a hash collision occurs.

It could also be worthwhile to translate the entire transition instance analysis algorithm described in Section 7.1 to one C function for each transition. Then the unification stack, a dynamic data structure, could be eliminated and replaced with a set of nested loops, each iterating the bindings for one symbolic token.

Chapter 9

Modelling Computer Programs

One of the major goals of this research is to provide a mechanism for semi-automatic verification of concurrent computer software, such as implementations of communications protocols. Presenting a complete set of rules for translating concurrent programs into algebraic system nets is out of the scope of this work, but the data type system has been created with this transformation in mind. We begin by summarising the features of our data type system and comparing them with other approaches. Later on, we sketch the transformation rules for some typical constructs and make some comparisons to related work.

9.1 Data Types

The rôle of powerful data types may not be ignored. While it is theoretically possible to represent all data by using integer numbers, it is very difficult to translate expressions (algebraic terms) manipulating structured data to equivalent expressions that work with the integer representations of the data objects. This has been tried in [33, 42], where problems were encountered e.g. with array indices. Another way to handle data types is to restrict the source language to allow only easily representable data types [20], severely restricting the usability of the verification tool.

Table 9.1 summarises the data types supported by some reachability analysers. A solid circle (●) denotes that the analyser supports the data type, whereas a hollow circle (○) means that a feature is partially supported or there is a mapping to a native data type of the analyser, but not all constraints are enforced or operations are supported by the analyser.

The four tools we compare to our reachability analyser for Algebraic System Nets, MARIA [41], are PROD [62], an analyser for Predicate/Transition Nets [14], DESIGN/CPN [43], a modelling tool for Coloured Petri Nets [31], and a protocol analyser SPIN [24], which is based on its own formalism, concurrent processes communicating via message queues.

Of the four reachability analysers, PROD has the scarcest data type system. Its data model is based on tuples of nonnegative integer numbers, which may be assigned lower and upper limits. The constraints in MARIA and DESIGN/CPN

Table 9.1: Data Types Representable in Different Reachability Analysers

	PROD	SPIN	MARIA	DESIGN/CPN
Constraints	○		●	●
Simple Data Types				
Boolean	○	●	●	●
Character	○	●	●	●
Integer	●	●	●	●
Enumerated	●	○	●	●
Identifier	○	○	●	○
Floating-Point Number				●
Character String		○	○	●
Structured Data Types				
Tuple	●	●	●	●
Array		●	●	●
Tagged Union	○	○	●	○
Variable-Length Buffer		○	●	●
Linked List				●

are based on Boolean conditions; it is possible to define types with “holes” in them, e.g. $\{i \in \mathbb{N} \mid i \leq 42 \vee 64 \leq i \leq 128\}$. There is not much variation among the simple data types. DESIGN/CPN, whose algebraic model is based on the Standard Meta Language [44], has two simple data types that fundamentally differ from the integer-like types: unbounded strings of characters and floating-point numbers. Character strings can be mapped to fixed-length character arrays in MARIA and SPIN. Floating point arithmetics is not supported in the other reachability analysers, probably because of a deliberate design decision. Floating point arithmetics is highly implementation-dependent, and results obtained involving it depend on the underlying computer architecture.

With structured data types, there are more differences between the analysers. All analysers support the tuple type. PROD does not support nested tuples, but this does not limit the expressive power of its type system, since there are no other structured data types in PROD. SPIN has fixed-length arrays indexed by integers, and the linked lists of DESIGN/CPN can be constrained to be of a constant length. In MARIA, arrays can be declared with any data type for indices, but one should keep in mind that an array with n different elements and m distinct indices has n^m possible values. Arrays indexed by ordered non-integer types can be represented with integer-indexed arrays by defining an enumerative mapping for the index values; an example can be found in Section 5.1.1.

Only MARIA supports the tagged union data type, which is necessary for modelling object-oriented constructs (Section 9.5) and for supporting the data type systems of some programming languages. Of the other analysers, PROD, whose data type system is limited to tuples of integers, can rather efficiently represent tagged unions of its natively supported data types. The mapping is straightfor-

ward: prepend each tuple with a tag element. In the other tools, one can form a tuple of an integer and all the data types belonging to the tagged union. The integer will indicate the active union component. This is utterly inefficient: while a tagged union of data items with domain sizes n_i for $1 \leq i \leq k$ has a domain of $\sum_{i=1}^k n_i$ elements, the domain size of the corresponding tuple mapping is $k \prod_{i=1}^k n_i$.

The linked list data type of DESIGN/CPN, like the unbounded character string, has an unbounded domain. When a list is constrained to have a maximum size, it is equivalent to the variable-length buffer of MARIA. SPIN has a separate language construct for defining message channels (queues with bounded capacity); the channels are not part of the data type system.

9.1.1 Expressive Power

Of the data type systems compared in Table 9.1, the data type system in PROD has the least expressive power. The SPIN data type system appears to be a proper subset of the system implemented in MARIA, and all data types of MARIA can be mapped to DESIGN/CPN data types.

9.1.2 Representation

The fundamental difference between the data type systems supported by MARIA and DESIGN/CPN is the boundedness of data types. Bounded systems have one disadvantage—possible lack of expressive power—and several advantages.

When a data type has a domain of known size n , each value can be represented with $\lceil \log_2 n \rceil$ bits. This affects especially the representation of the state space. The encoding described in Section 7.3 clearly outperforms the one used in DESIGN/CPN, which maintains the data in the pointer structures of the underlying Standard Meta Language [44] implementation.

The data types in PROD have bounded domains, but the analyser does not encourage the use of constraints. Therefore, the state space encoder often has to deal with tuples of unconstrained machine words. It assumes small tuple elements to be more frequent than larger elements, and encodes them with fewer bits.

9.2 Message Queues

The Open Systems Interconnection Reference Model [26] and related models provide a common basis for communications protocols. Protocol entities in open systems communicate with each other by sending messages via lower-level entities, which are connected via a physical medium to the underlying network. The messages are usually buffered at the receiving end obeying the first-in-first-out discipline. For instance, the dynamic semantics of SDL, the CCITT Specification and Description Language [30], builds heavily on such buffering of messages, which are called “signals” in SDL.

In the idealistic world of SDL, buffers have infinite capacity, while all practical computing systems have a finite amount of memory. In many applications, “astro-

nomically large” is a good substitute for infinite. When applying formal methods, especially in exhaustive reachability analysis, even models with buffer capacities limited to four messages may be unanalysable, and one has to hope that if there are errors in the protocol, they will be discovered also with severely limited capacities.

9.2.1 Previous Approaches

It is possible to represent message queues without introducing a data type for them. If a queue can hold only one kind of data items or if the semantics of the “queue” allows enqueued items to be fetched in arbitrary order, the queue can be represented with one simple net place. Usually this is not the case.

We mention two previous approaches for modelling queues. Neither of them maps message queues directly to native data types, which would atomise queue operations.

Managing Queues as Circular Arrays

The approach taken in [33, 42] essentially models the queues of the SDL variant TNSDL [52] as circular arrays. It uses one place containing tokens (items in the buffer) tagged with the buffer position, and other places holding pointers to the front and back of the queue.

The advantage of this approach is that the queue contents never will be shifted. When an item is removed from the buffer, its slot will be marked empty, and the pointer to the front of the queue will be updated.

Unfortunately, the advantage is also a disadvantage. For instance, a first-in-first-out buffer of capacity 4 containing the items a , b and c could be represented as $\langle\langle a, b, c, \varepsilon \rangle, \langle 0, 3 \rangle\rangle$ but also as $\langle\langle c, \varepsilon, a, b \rangle, \langle 2, 1 \rangle\rangle$. The two structures are identical only when interpreted in an appropriate way. A simple reachability analyser, which does not see the similarity, may generate many uninteresting “copies” of the essential state space of a system modelled in such a way.

Managing Queues by Shifting Items

A similar approach, represented in [20], avoids the problem of seemingly different states by not using pointers. Instead, the buffer contents will be shifted, one item at a time, when an item is read from the buffer. The shifting causes numerous uninteresting intermediate states. Especially if a model encompasses several buffers of this kind that can be used concurrently, its reachability graph is likely to contain several lattices resembling those depicted in Section 2.2.

9.2.2 Algebraic Support for Queues

The variable-length buffer data type described in Section 5.3.3 suits to modelling message queues. The contents of the message queue can be packed in a variable whose type encapsulates a buffer over the message type. More precisely: assuming

that some $s \in \mathcal{S}_\beta$ is the message sort and $\mathcal{D}_s^{\mathcal{A}}$ its carrier, the carrier of the associated message buffer sort s' could be $\mathcal{D}_{s'}^{\mathcal{A}} = \mathbb{V}_{\mathcal{D}_s^{\mathcal{A}}, n}$ for some buffer length n .

The carrier of s' could also be some other data type encapsulating $\mathbb{V}_{\mathcal{D}_s^{\mathcal{A}}, n}$. For instance, if the system consists of a number of similar buffers, it could make sense to represent the buffers as an array of buffers, or as sets of tuples with some components identifying the “owner” of the buffer and one component holding the buffer contents.

When the underlying formalism includes basic queue operations, there is no need to pay attention to the shifting operations discussed earlier. All algebraic terms in our class of Algebraic System Nets are evaluated atomically; this includes all queue operations.

Some languages include operations that break the first-in-first-out principle of queues. For instance, the *save* operation of SDL [30] makes it possible to skip over certain messages in the front of a queue and to process and dequeue a following message. In order to facilitate simple transformations for this kind of non-orthodox operations, it is advisable to foresee the modelling formalism with an algebraic operation having enough expressive power. The queue operations defined in Section 6.3.5 include indexed variants of the basic operations.

9.3 Dynamic Resource Allocation

It is straightforward to model systems that use a fixed amount of resources over their lifetime. For systems that allocate resources dynamically, it is often difficult to foresee the maximum number of resources it will use in any of its reachable states. In practice, systems always have some limited number of resources, and an allocation operation may fail due to lack of available resources. If the pool of available resources is big enough so that a resource shortage never occurs, the illusion of a truly dynamic allocation of resources from an infinite pool works.

So, resource allocation can be modelled with some finite pool. In such a model, there is a fixed amount of resources, some of which are available for allocation. Similar to limiting the capacity of message queues (Section 9.2), the resource pool usually has to be kept small in order to be able to analyse the model.

Although dynamic resource management is not particularly difficult to model, it has been neglected in many approaches. For instance, the transformation presented in [20] does not allow dynamic process creation.

9.3.1 Process Creation

The computations of a concurrent system are often executed in processes, each one executing its own algorithm, synchronising with other processes from time to time. Some systems are implemented with a dynamic number of processes. For instance, a server might have one master process that listens for requests and creates slave processes, each of which serves one request before terminating.

The dynamic resource involved are data structures. Processes have a local state, which has to be maintained somehow, and in order to allow inter-process

communication, each process has to be given a process identifier. A straightforward approach of modelling processes with Algebraic System Nets is to introduce a data type for the process identifier and to store the local variables of each process into data structures that can be indexed by the process identifier, e.g. into places whose carrier consists of multi-sets over tuples, one of whose components is the process identifier.

The identifier type defined in Section 5.2.5 suits very well to describing process identifiers. A pool of available process identifiers can be modelled with a place that initially contains all process identifiers. Each process creation operation will consume an identifier from the pool, and each process termination will restore the identifier to the pool.

When a process terminates, it may be useful to verify that no other processes hold the process identifier. Let us assume that there are processes “thinking” that a particular process still exists, even though it has terminated. When the process terminated, it returned its identifier back to the pool, and a new process creation operation may assign the same identifier to some other process. The processes still holding the old identifier may then mistake the new process for the old one.

This kind of errors, *dangling resource identifiers*, can presently be detected with model checking, but the associated temporal logic formulae or property automata may become very large. It would be tempting to introduce an algebraic operation that maps a marking of a net and an identifier value to a truth value:

$$J : \mathcal{M}_\Sigma \times \mathbb{J}_n \rightarrow \mathbb{B} : \langle M, j \rangle \mapsto \begin{cases} \top & \text{if “} j \notin M(p) \text{” for all } p \in P \\ \perp & \text{otherwise.} \end{cases}$$

The informal notation “ $j \notin M(p)$ ” denotes that the identifier value j does not occur anywhere in the place marking, not even in a component of a structured value.

This topic deserves further research.

9.3.2 Memory Allocation

As discussed in Section 3.2.1, data structures should be abstracted away from formal models. Industrial-size programs exhibiting memory management problems typically are too big to be analysed using formal methods. One must try to cope with the less complete methods described in Section 3.1, such as static analysis and performing regression testing on instrumented program code.

But what if one absolutely must create a model involving dynamic memory allocation? Like with dynamic process creation, the answer lies in the identifier data type, which can represent pointers to data items. When a program needs to allocate data items of sort s whose carrier is \mathcal{D}_s^A , the generated model will contain an identifier pool (sort s' , carrier $\mathcal{D}_{s'}^A$) for the “pointers” of \mathcal{D}_s^A , and the dynamically allocated data items of sort s will be represented using a sort s'' having the carrier $\mathcal{D}_{s''}^A = \mathbb{T}_{\mathcal{D}_{s'}^A, \mathcal{D}_s^A}$.

This approach uses disjoint sets of pointers for different data types. As always with the identifier type, one has to assign a maximum number of simultaneously allocated data items for each dynamically allocated data type. The phenomenon

with dangling resource identifiers concerns also memory allocation, since it is possible to deallocate a memory object to which pointers still exist in data structures.

With this approach, it is possible to model dynamic memory management that relies on explicit deallocation of unused data items. Systems that perform *garbage collection*, automatically deallocating data items which are no longer used, are difficult to model in an efficient way. One option would be to extend the formalism with a garbage collector, but an extension of such a global nature would affect many analysis methods and algorithms.

9.4 Procedure Calls

Structured programming languages require that program statements are encapsulated into units referred to as procedures. Procedures are very useful for structuring program code, and they can hide obscuring details of low-level operations from high-level procedure code. An Algebraic System Net, however, is a flat formalism: there are just places and transitions, all on the same level. In order to translate a structural program to an Algebraic System Net, the structure must be *flattened*, losing the structural information.¹

Typical programming languages define procedure calls as part of the expression syntax. In other words, the underlying many-sorted algebra of a programming language may contain procedure calls as terms. This provides a convenient way of defining operations without extending the language core.

Allowing procedure calls in algebraic terms poses a problem when the terms are to be translated to a less expressive algebra that does not encapsulate procedure calls. Only calls to simple procedures that do not modify their environment e.g. by altering some data structures or by sending messages over the network, can be translated directly.

One solution is to divide the terms to subterms not containing procedure calls, and to evaluate the expression in several steps, making the procedure calls in appropriate places. In this way, one statement in a program can be translated into several transitions in an Algebraic System Net. This approach is analogous with the one outlined in Section 7.2.3. The transformation given in [20] transforms all procedure calls to several transitions, without giving special treatment to purely functional procedures that do not modify their environment. In [33, 42], procedure calls are treated as statements, so they cannot occur in subexpressions.

9.4.1 Scoping

Procedures typically have their own *name spaces*: variable x in procedure a is different from variable x in procedure b . It is thus impossible to represent both variables with just one place \mathbf{x} that would hold the values of the variables. An

¹When generating a model of a program, one can provide some structural information by assigning meaningful identifiers to places and transitions. For instance, transition identifiers could consist of the name of the program file and the line and column numbers where the corresponding statement or subexpression of the program begins.

obvious solution is to incorporate the *scope* identifier to the label, creating e.g. places **A:X** and **B:X**.

9.4.2 Recursive Procedures

Non-recursive procedure calls can be translated with a simple macro expansion, replacing each procedure call with the body of the called procedure, substituting the procedure parameters. The transformation represented in [20] is equivalent to a macro expansion in its expressive power: it cannot handle cyclic or recursive procedure invocations.

Recursive procedures can easily be implemented with a stack architecture. Each procedure invocation is executed in its own context. The contexts can be identified e.g. with the current recursion depth, the current number of unfinished procedure invocations. The *control flow* of the procedures, the order in which the statements will be executed, can be modelled with *control places*. Each statement will be translated to a sequence of transitions moving a *control token* from one control place to another. When several invocations of a procedure can be active at the same time, the control token of the procedure must contain the context identifier, e.g. the recursion depth.

Since the data types in our class of Algebraic System Nets are finite, we cannot allow infinite recursion depth. Analogous to dynamic resource management, a limit must be fixed for the recursion depth. If it is big enough and if the system being analysed does not exhibit infinite behaviour, the limit will not restrict the behaviour of the system.

9.4.3 Exception Handling

Exception handling [18] affects the semantics of procedure calls. When an unexpected condition occurs during the execution of a procedure, the procedure or the run-time system may raise an exception, transferring control to the nearest applicable exception handler. If no exception handler can be applied in the procedure, the exception will be raised further in the calling procedures, until an applicable exception handler is found.

The framework proposed in [18] allows resumption of the action that raised the exception. Most modern languages do not allow that: it is impossible for an exception handler to return to the statement that raised the exception condition. This means that exceptions can be modelled as a special kind of jump operations, which can return from several levels of procedure calls in one step, similar to the `setjmp` and `longjmp` operations defined in the C programming language [28].

9.5 Object-Oriented Constructs

The concept of object-oriented programming, the buzzword of the past decade, was introduced in the early 1960s. One of the first implementations, and the ori-

ginal inspiration behind C++ [27] was Simula [7], a language designed for writing simulations.

Object-oriented programming ties data type definitions and operations syntactically together. Traditional programming languages based on many-sorted algebras have typed (sorted) variables and operations whose signatures indicate the argument types. Object-oriented programming languages have a special kind of data types, called *classes*, which have *member variables* and *methods*. A class definition can be seen as a tuple data type definition, with the member variables as the components of the tuple. Methods are algebraic operations whose definition is syntactically bundled with the class definition. The instances of a class, or the variables of a class sort, are called *objects*.

The first C++ implementation [56] translated the object-oriented constructs to C code, mapping classes (object types) to structured data types, and rewriting identifiers to be compatible with the flat name space of C. We suggest a similar approach, flattening the constructs, for translating object-oriented programs to Algebraic System Nets.

9.5.1 Inheritance and Polymorphism

Simple objects, class instances, can be represented as tuples of their member variables. This is not the whole truth. Objects can *inherit* properties from each other. A class *derived* from a *base class* inherits all the member variables of the base class, and it may add new member variables. There can be inheritance on multiple levels, and several classes can derive from the same base class. Nevertheless, the inheritance structure can be represented as a tree or as a directed acyclic graph, if *multiple inheritance*, deriving properties from more than one base class, is allowed.

Objects of derived classes, which can sometimes also be used as objects of a base class, can be represented with a union of structured data types. Consider a base class C whose members can be represented with the carrier \mathcal{D}_s^A of a sort s . A derived class C' whose own members can be represented with the carrier $\mathcal{D}_{s'}^A$ of another sort s' , can be represented as the tuple $\mathbb{T}_{\mathcal{D}_s^A, \mathcal{D}_{s'}^A}$.

It is possible for a derived class to *override* some of the methods defined for the base class. When a particular method is applied to an object, the corresponding method defined for the base class will only be used if it has not been overridden in a derived class whose instance the object is. A method that can do different things on different kinds of objects is called *polymorphic*.

Polymorphism can be realised by allowing objects to belong either to a base class or to one of its derived classes. Such objects can be represented as a tuple of the carrier of the base class and a tagged union over the carriers of the derived class members, and over a singleton type. The singleton type applies for objects of the base class, which have no derived members. The algebraic operations for the union type defined in Sections 6.3.5 and 6.3.6 can be used to find out the classes whose instance an object is, and to convert objects of derived classes to objects of corresponding base classes.

Chapter 10

Conclusion

10.1 Contribution of the Thesis

There exist numerous reachability analysers for models of concurrent and distributed systems, but most of them lack a powerful data type system and algebraic operations that would facilitate a straightforward transformation from a computer program to a formal analysis model. Such tools are typically used for analysing manually constructed models.

The tedious and time-consuming task of constructing a formal verification model of a system can be eased by developing more powerful tools. We have implemented a freely distributable computer tool for analysing models employing a data type system that tolerates the comparison with the data type systems used in some state-of-the-art reachability analysers.

While the formalism we deploy has not specifically been designed for constructing models of computer software, the transformation principles outlined in Chapter 9 should cover the intrinsic constructs of commonly used procedural and object-oriented programming languages. Some of the transformations have been presented in earlier work; some of them appear to be new.

Our description formalism, Algebraic System Nets [37], has a solid mathematical foundation, and there is active research on efficient analysis methods for the formalism. A substantial part of the research concentrates on a lower-level formalism, Place/Transition Nets [50]. Analysis methods developed for that level can be lifted to our class of nets, since all models constructed in our formalism can be unfolded to the lower level.

The framework of our formalism is based on many-sorted algebras, which are defined in Chapter 4. The interpretation rules of our many-sorted algebras include short-circuit evaluation, a common feature in programming languages that is easily overlooked. Short-circuit evaluation is not merely an optimisation; it also affects the semantics when a subexpression skipped due to short-circuit evaluation could not be evaluated due to an error.

Implementing an expression evaluator is a rather irksome task, but there is some place for important design choices, such as different optimisations, or the way errors are reported. Especially the latter greatly affects the usability of the expression evaluator, both from the user and from the programmer perspective.

The biggest user of an expression evaluator in a Petri Net reachability analyser is the transition instance analysis algorithm described in Section 7.1, which finds all actions that can be performed in a given system state. If the instance analyser silently ignored erroneous actions, as at least one implementation [62] does, serious errors could be overlooked, or they could be hard to locate. Our instance analyser reports all erroneous transition instances, even incomplete ones, and the user will always be notified when something is wrong in the model.

In addition to the common data types and basic algebraic operations, our analyser tool implementation [41] defines a total order and a compact encoding for all data types, including structured and constrained data types.

Compact encoding is essential for space efficient management of the reachability graphs of complex systems. In our formalism, the nodes of a reachability graph are sequences of multi-sets over sorted values. We believe the encoding method described in Section 7.3.1 to be an original invention. According to Table 8.1, it performs an order of magnitude better than the method used in a comparable tool.

10.2 Future Work

Currently the reachability analyser we have designed and implemented lacks some essential features, such as on-the-fly model checking of properties specified in temporal logic with the presence of fairness constraints, and various reduction methods that cause uninteresting intermediate states to be omitted from the reachability graph. These areas are being investigated, and the implementation is being worked on.

This work discusses language-specific front-ends, which translate a system description given in an application-specific high-level language to a formal model that can be analysed, but does not refer to an actual implementation. Our compiler [40] for SDL [30] is being extended with model generation routines, so that specifications of telecommunications protocols can be input to the reachability analyser. It will be an interesting exercise to implement transformations for true object-oriented constructs and exception handling, which should be included in the upcoming version of SDL.

Bibliography

- [1] Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers—Principles, Techniques, and Tools*. Addison–Wesley, Reading, MA, USA, March 1986.
- [2] Brian Berliner. CVS II: Parallelizing software development. In *Winter 1990 USENIX Conference*, pages 341–352, Washington, DC, USA, 1990.
- [3] Jonathan Billington. Many-sorted high level nets. In *3rd Workshop on Petri Nets and Performance Models*, pages 166–179, Washington, DC, USA, 11–13 December 1989. IEEE CS Press, Los Alamitos, CA, USA.
- [4] Edmund M. Clarke Jr., Orna Grümberg, and Doron A. Peled. *Model Checking*. MIT Press, Cambridge, MA, USA, 1999.
- [5] Danny Cohen. On holy wars and a plea for peace. *IEEE Computer*, 14(10):48–54, October 1981.
- [6] Robert Corbett, Richard Stallman, and Wilfred Hansen. *Bison Manual: Using the YACC-compatible Parser Generator, for Version 1.29*. Free Software Foundation, Boston, MA, USA, November 1999.
- [7] Ole-Johan Dahl and Kristen Nygaard. SIMULA—an Algol based simulation language. *Communications of the ACM*, 9(9):671–678, September 1966.
- [8] Digital Equipment Corporation, Maynard, MA, USA. *Using DIGITAL C++ for DIGITAL UNIX Systems*, December 1997. Software Version 6.0.
- [9] Jeremy Dion and Louis Monier. *Third Degree: heap usage and leak profiler, and memory-access error checker for C and C++ programs*. Digital Equipment Corporation, Maynard, MA, USA, 1998.
- [10] Anders Ek. Automatic debugging of communicating systems using the SDT Validator. Technical paper, Telelogic AB, Malmö, Sweden, September 1998.
- [11] Frits Feldbrugge. List of Petri Net tools. *Petri Net Newsletter*, 22:20–32, October 1985.
- [12] Free Software Foundation, Boston, MA, USA. *GNU Emacs Manual, for Version 20.1*, 1997.

- [13] Jaco Geldenhuys and Pieter de Villiers. Runtime efficient state compaction in SPIN. In Dennis Dams and Mieke Massink, editors, *The 5th International SPIN Workshop on Theoretical Aspects of Model Checking*, pages 3–12, Trento, Italy, July 5 1999.
- [14] Hartmann J. Genrich. Predicate/Transition Nets. In Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Petri Nets: Central Models and their Properties—Advances in Petri Nets 1986, Part I, Proceedings of an Advanced Course*, volume 254 of *Lecture Notes in Computer Science*, pages 207–247, Bad Honnef, Germany, September 1986. Springer-Verlag, Berlin, Germany, 1987.
- [15] Tristan Gingold. *Checker, a memory access detector*, 1996. Documentation for Version 0.8.
- [16] Patrice Godefroid and Gerard J. Holzmann. On the verification of temporal properties. In *13th IFIP Symposium on Protocol Specification, Testing and Verification*, pages 109–124, Liège, Belgium, May 1993.
- [17] Joseph A. Goguen. Abstract errors for abstract data types. In Erich J. Neuhold, editor, *Formal Description of Programming Concepts*, pages 491–525. North-Holland Publishing Company, Amsterdam, The Netherlands, 1978.
- [18] John B. Goodenough. Exception handling: Issues and a proposed notation. *Communications of the ACM*, 18(12), December 1975.
- [19] Bernd Grahlmann. The PEP tool. In Orna Grumberg, editor, *Computer Aided Verification, 9th International Conference, CAV '97*, volume 1254 of *Lecture Notes in Computer Science*, pages 440–443, Haifa, Israel, June 1997. Springer-Verlag, Berlin, Germany.
- [20] Bernd Grahlmann. *Parallel Programs as Petri Nets*. Dissertation, Universität Hildesheim, Fachbereich Mathematik, Informatik, Naturwissenschaften, Hildesheim, Germany, September 1998.
- [21] Jean-Charles Grégoire. State space compression in SPIN with GETSs. In *2nd International SPIN Verification Workshop*, New Brunswick, NJ, USA, August 1996.
- [22] Keijo Heljanko. Model checking the branching time temporal logic CTL. Research Report A45, Helsinki University of Technology, Department of Computer Science and Engineering, Digital Systems Laboratory, Espoo, Finland, May 1997.
- [23] Hewlett-Packard Company, Palo Alto, CA, USA. *HP aC++ Version A.03.10 Release Notes*, June 1998.
- [24] Gerard J. Holzmann. *Design and Validation of Computer Protocols*. Prentice Hall, Englewood Cliffs, NJ, USA, 1991.

- [25] Nisse Husberg. Verifying SDL programs using Petri nets. In *1998 IEEE International Conference on Systems, Man, and Cybernetics*, volume 1, pages 208–213, San Diego, CA, USA, October 1998. Institute of Electrical and Electronics Engineers, Inc.
- [26] *Information Processing Systems—OSI Reference Model—The Basic Model*. ISO/IEC 7498-1. International Organization for Standardization, Geneva, Switzerland, 1994.
- [27] *Information Technology—Programming Languages—C++*. ISO/IEC 14882. International Organization for Standardization, Geneva, Switzerland, 1998.
- [28] *Information Technology—Programming Languages—C*. ISO/IEC 9899. International Organization for Standardization, Geneva, Switzerland, 1999.
- [29] *Information Technology—Portable Operating System Interface (POSIX)—Part 2: Shell and Utilities*. ISO/IEC 9945-2. International Organization for Standardization, Geneva, Switzerland, 1993.
- [30] *CCITT Specification and Description Language (SDL)*. Recommendation Z.100. International Telecommunication Union, Geneva, Switzerland, October 1996.
- [31] Kurt Jensen. *Coloured Petri Nets: Basic Concepts, Analysis Methods and Practical Use: Volume 2, Analysis Methods*. Monographs in Theoretical Computer Science. Springer-Verlag, Berlin, Germany, 1995.
- [32] Tommi Junttila. *Detecting and Exploiting Data Type Symmetries of Algebraic System Nets during Reachability Analysis*. Licentiate's thesis, Helsinki University of Technology, Department of Computer Science and Engineering, Espoo, Finland, December 1999.
- [33] Tero Jyrinki. Dynamic analysis of SDL programs with Predicate/Transition Nets. Technical Report B17, Helsinki University of Technology, Department of Computer Science and Engineering, Digital Systems Laboratory, Espoo, Finland, April 1997.
- [34] Jukka Kemppainen. *ARA Tools, Reference Manual*. VTT Electronics, Oulu, Finland, March 1994.
- [35] Esa Kettunen, Esa Montonen, and Timo Tuuliniemi. An interactive PrT-net tool for verification of SDL-specifications. Technical Report B3, Helsinki University of Technology, Department of Computer Science and Engineering, Digital Systems Laboratory, Espoo, Finland, February 1988.
- [36] Ekkart Kindler and Wolfgang Reisig. Algebraic System Nets for modelling distributed algorithms. *Petri Net Newsletter*, 51:16–31, December 1996.

- [37] Ekkart Kindler and Hagen Völzer. Flexibility in algebraic nets. In Jörg Desel and Manuel Silva, editors, *Application and Theory of Petri Nets 1998: 19th International Conference, ICATPN'98*, volume 1420 of *Lecture Notes in Computer Science*, pages 345–364, Lisbon, Portugal, June 1998. Springer-Verlag, Berlin, Germany.
- [38] Raimo Kujansuu, Leo Ojala, and Heikki Tuominen. Verification and validation of communication protocols. In Carl A. Sunshine, editor, *Protocol Testing, Specification and Validation*, pages 311–313, Idyllwild, CA, USA, May 1982. North-Holland Publishing Company, Amsterdam, The Netherlands.
- [39] Marek Leszak and Horst Eggert. *Petri-Netz-Methoden und -Werkzeuge*, volume 197 of *Informatik-Fachberichte*. Springer-Verlag, Berlin, Germany, October 1988.
- [40] Marko Mäkelä. Implementing the front-end of an SDL compiler. Master's thesis, Helsinki University of Technology, Department of Computer Science and Engineering, Espoo, Finland, December 1998.
- [41] Marko Mäkelä. *Maria—Modular Reachability Analyzer for Algebraic System Nets*. Helsinki University of Technology, Laboratory for Theoretical Computer Science, Espoo, Finland, January 2000. On-line documentation, [URL: `http://www.tcs.hut.fi/maria/`](http://www.tcs.hut.fi/maria/).
- [42] Markus Malmqvist. Methodology of dynamical analysis of SDL programs using Predicate/Transition Nets. Technical Report B16, Helsinki University of Technology, Department of Computer Science and Engineering, Digital Systems Laboratory, Espoo, Finland, April 1997.
- [43] Meta Software Corporation, Cambridge, MA, USA. *Design/CPN Reference Manual for X-Windows, Version 2.0*, 1993.
- [44] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML*. MIT Press, Cambridge, MA, USA, 1990.
- [45] Daniel P. Miranker. TREAT: A better match algorithm for AI production systems. In Howard Forbus and Kenneth Shrobe, editors, *Proceedings of the 6th National Conference on Artificial Intelligence*, pages 42–47, Seattle, WA, USA, July 1987. Morgan Kaufmann, San Mateo, CA, USA.
- [46] Christos H. Papadimitriou. *Computational Complexity*. Addison–Wesley, Reading, MA, USA, 1994.
- [47] Vern Paxson, Van Jacobson, Jef Poskanzer, and Kevin Gong. *Flex: The Lexical Scanner Generator, for Version 2.5*. Free Software Foundation, Boston, MA, USA, April 1995.
- [48] Bruce Perens. *Electric Fence Malloc Debugger*. Pixar Animation Studios, 1993. Manual page for Version 2.0.5.

- [49] Carl Adam Petri. *Kommunikation mit Automaten*. Dissertation, Technische Universität Darmstadt, Fachbereich Mathematik, Physik, Darmstadt, Germany, 1962.
- [50] Wolfgang Reisig. *Petrinetze—Eine Einführung*. Springer-Verlag, Berlin, Germany, 1986.
- [51] Wolfgang Reisig. Petri nets and algebraic specifications. *Theoretical Computer Science*, 80:1–34, March 1991.
- [52] Erkki Ruohtula, Esa Kettunen, and Heikki Tuominen. *TNSDL Book*. Nokia Telecommunications, Inc., 3rd edition, November 1995.
- [53] Sriram Srinivasan. *Advanced Perl Programming*. O'Reilly & Associates, Sebastopol, CA, USA, 1st edition, August 1997.
- [54] Richard M. Stallman. *Using and Porting GNU CC*. Free Software Foundation, Boston, MA, USA, 1999. Software Version 2.95.
- [55] Harald Störrle. An evaluation of high-end tools for Petri-nets. Bericht 9802, Ludwig-Maximilians-Universität München, Institut für Informatik, Munich, Germany, June 1998.
- [56] Bjarne Stroustrup. Adding classes to the C language: An exercise in language evolution. *Software—Practice and Experience*, pages 139–161, February 1983.
- [57] Walter F. Tichy. RCS: A system for version control. *Software—Practice and Experience*, 15(7):637–654, July 1985.
- [58] Linus Torvalds. Linux: a portable operating system. Master's thesis C-1997-12, University of Helsinki, Department of Computer Science, Helsinki, Finland, 1997.
- [59] Antti Valmari. State space generation: Efficiency and practicality. Publications 55, Tampere University of Technology, Tampere, Finland, 1988.
- [60] Antti Valmari. The state explosion problem. In Wolfgang Reisig and Grzegorz Rozenberg, editors, *Lectures on Petri Nets I: Basic models*, volume 1491 of *Lecture Notes in Computer Science*, pages 429–528. Springer-Verlag, Berlin, Germany, 1998.
- [61] Kimmo Varpaaniemi. On the stubborn set method in reduced state space generation. Research Report A51, Helsinki University of Technology, Department of Computer Science and Engineering, Digital Systems Laboratory, Espoo, Finland, May 1998.
- [62] Kimmo Varpaaniemi, Jaakko Halme, Kari Hiekkänen, and Tino Pyssysalo. PROD reference manual. Technical Report B13, Helsinki University of Technology, Department of Computer Science and Engineering, Digital Systems Laboratory, Espoo, Finland, August 1995.

- [63] Larry Wall, Tom Christiansen, and Randal L. Schwartz. *Programming Perl*. O'Reilly & Associates, Sebastopol, CA, USA, 2nd edition, May 1997.
- [64] Gary Watson. *Debug Malloc Library*, January 1998. Documentation for Version 3.3.1.
- [65] Colin H. West. Automated protocol validation. In Carl A. Sunshine, editor, *Protocol Testing, Specification and Validation*, pages 361–371, Idyllwild, CA, USA, May 1982. North-Holland Publishing Company, Amsterdam, The Netherlands.

Index

- abstractions, *see* models, making \sim
- algebraic definitions
 - \approx_v (term compatibility), 51
 - \mathcal{A} (algebra), 22
 - A_v (assignment candidates), 50
 - $\mathcal{D}^{\mathcal{A}}, \mathcal{D}_s^{\mathcal{A}}$ (carriers), 22
 - $\check{\mathcal{D}}^{\mathcal{A}}, \check{\mathcal{D}}_s^{\mathcal{A}}$ ($\mathcal{D}^{\mathcal{A}} \cup \{\varepsilon\}$), 22
 - ε (undefined symbol), 22
 - $e_v^{\mathcal{A}}(T)$ (evaluation of term T), 22
 - \mathcal{F} (function symbols), 21
 - $\mathcal{F}^{\mathcal{A}}$ (operations), 22
 - $\mathcal{F} \subseteq (\mathcal{T} \times \mathcal{P}) \cup (\mathcal{P} \times \mathcal{T})$ (flow relation), 25
 - \mathcal{G} (short-circuit symbols), 21
 - $i: \mathcal{P} \cup \mathcal{T} \cup \mathcal{F} \rightarrow \mathbf{T}^{\mathbf{S}\mu}$ (net inscriptions), 25
 - $M \in \mathcal{M}_{\Sigma}$ (marking), 26
 - $M_0 \in \mathcal{M}_{\Sigma}$ (initial marking), 26
 - $\mu \in \mathcal{M}(A)$ (multi-set over A), 23
 - $\mathcal{N} = \langle \mathcal{P}, \mathcal{T}, \mathcal{F} \rangle$ (net), 25
 - \mathcal{P} (places), 25
 - $\Sigma = \langle \mathcal{N}, \mathcal{A}, \mathcal{V}, i \rangle$ (algebraic system net), 25
 - \mathbf{S} (signature), 21
 - \mathbf{S}_{μ} (multi-set signature), 24
 - \mathcal{S} (sorts), 21
 - \mathcal{S}_{β} (basic sorts), 24
 - \mathcal{S}_{μ} (multi-set sorts), 24
 - $\mu: \mathcal{S}_{\beta} \rightarrow \mathcal{S}_{\mu}$ (sort mapping), 24
 - $\mathbf{T}^{\mathbf{S}}(\mathcal{V}), \mathbf{T}_s^{\mathbf{S}}(\mathcal{V})$ (terms), 21
 - \mathcal{T} (transitions), 25
 - t_v^+ (output effect), 26
 - t_v^- (input effect), 26
 - $\mathcal{V}, \mathcal{V}_s$ (variables), 21
 - $V^{\mathcal{A}}(\mathcal{V})$ (assignments over \mathcal{V}), 22
 - $v \in V^{\mathcal{A}}(\mathcal{V})$ (an assignment), 22
- algebraic operations, 22, 36–48
 - and**, 39
 - bitand**, 40
 - bitnot**, 40
 - bitor**, 40
 - bitxor**, 40
 - card** _{s} , 47
 - comparison, 38, 47
 - component** _{s,k} , 41, 42
 - cons** _{s} , 41
 - cons** _{s,k} , 42
 - constant** _{d} , 37
 - convert** _{s,s'} , 43
 - defined** _{s,k} , 42
 - divide**, 40
 - empty** _{s} , 45
 - enqueue** _{s} , 41
 - enqueue-at** _{s} , 41
 - equal**, 38
 - equalset** _{s} , 47
 - equiv**, 39
 - ext** _{s} , 44
 - filter** _{x,T} , 45
 - free** _{s} , 41
 - greater**, 38
 - greaterequal**, 38
 - impl**, 39
 - index** _{s} , 41
 - int** _{s} , 44
 - less**, 38
 - lessequal**, 38
 - map** _{x,T} , 46
 - minus** _{s} , 46
 - minus**, 40
 - mmap** _{x,y,U} , 46
 - modulus**, 40
 - mset** _{s} , 44
 - mul** _{s} , 47
 - negate**, 40
 - not**, 39

- on basic sorts, 37–44
- on multi-set sorts, 44–47
- or**, 39
- peek**_s, 41
- peek-at**_s, 41
- plus**, 40
- pred**, 38
- push**_s, 41
- push-at**_s, 41
- remove**_s, 41
- remove-at**_s, 41
- select**_{s',s''}, 47
- shiffl**, 40
- shiftr**, 40
- short-circuit, 47–48
- subset**_s, 47
- succ**, 38
- sum**_{x,T,U}, 45, 59
- times**, 40
- undefined**, 38
- unequal**, 38
- union**_s, 46
- used**_s, 41
- xor**, 39
- algebraic system nets, 24–28
 - actions, 26–28
 - performing, 27–28
 - analysing, *see* models, \sim
 - carriers, *see* data types
 - markings, 26
 - operations, *see* algebraic \sim
- algebras, 20–24
 - multi-set, 24
 - terms, 21
 - evaluating, 22–23, 56–61
- arc, 25
 - expressions, *see* inscriptions
 - effect, 26
- assignments, 22
- carriers, *see* data types
- component
 - of a tuple, 32
 - of an array, 33
- control flow, 87
- data types, 29–35
 - $<_{\mathcal{D}}$ (total order on \mathcal{D}), 29
 - \mathcal{D} (data type), 29
 - $o_{\mathcal{D}} : \mathcal{D} \rightarrow \{0, \dots, |\mathcal{D}| - 1\}$ (bijective mapping), 22, 30
 - constraints, 34–35
 - simple, 30–32
 - \mathbb{B} (boolean), 31
 - \mathbb{K} (character), 31
 - \mathbb{E}_N (enumerated), 31
 - \mathbb{J}_n (identifier), 32
 - \mathbb{I} (integer), 31
 - structured, 32–34
 - $\mathbb{A}_{\mathcal{D}_x, \mathcal{D}_e}$ (array), 33
 - $\mathbb{V}_{\mathcal{D}_e, n}$ (buffer), 33
 - $\mathbb{T}_{\mathcal{D}_1 \dots \mathcal{D}_n}$ (tuple), 32
 - $\mathbb{U}_{\mathcal{D}_1 \dots \mathcal{D}_n}$ (union), 34
- expressions, *see* algebras, terms
 - arc, 25
 - effect, 26
 - evaluating, 22–23, 56–61
 - initialisation, 25
- family (of sets), 20
- formal methods, 14–18
- garbage collection, *see* resources
- input arc, 25
 - input effect, 26
- inscriptions, 25
- language-specific front-end, 15–17
- marking, 26
 - initial, 26
- model checking, 17
- modelling, 80–88
 - message queues, 82–84
 - object-orientation, 87–88
 - procedure calls, 86–87
 - resource allocation, 84–86
- models
 - analysing, 17–18, 50–66
 - enabled actions, 50–56
 - managing states, 61–66
 - unfolding, 55
 - constructing, 15–17, 72–88

- automatically, 80–88
 - manually, 72–79
- making abstractions, 8–11
 - atomicity, 8
 - drawbacks, 11
 - nondeterminism, 9
- multi-set, 23–24
 - algebras, 24
 - signature, 24
- multiplicity, 23
- operation symbols, 21
- operations, *see* algebraic ~
 - on basic sorts, 37–44
 - on multi-set sorts, 44–47
 - on multi-sets, 23
- output arc, 25
 - output effect, 26
- places, 25
 - capacity constraints, 63
 - implicit, 65, 76
 - invariants, 73
- pointers, 30, 32, 37, 85
- programs, 7–8
 - analysing, 12–18
 - exception handling, 87
 - instrumenting, 13
 - object-oriented, 87–88
 - procedures, 86–87
 - statements, 7, 36
 - control flow, 87
 - encapsulating, 86
 - executing, 87
 - unreachable, 12
 - testing, 14
- reachability graph, *see* state space
- resources
 - detecting leakages, 13–14
 - modelling, 84–86
- short-circuit evaluation, 21, 47–48
- signature, 21
 - multi-set, 24
- simulation, 17
- sort, 21
 - basic and multi-set, 24
- state space, 26
 - and data, 16
 - explosion, 17
 - generating, 28
 - locality, 30
 - reducing, 8
 - capacity constraints, 63
 - implicit places, 65, 73, 76
 - storing, 61–66
 - transitions, 65
- statements, *see* programs, ~
- symbols, 21
- syntactic sugar, 36
- terms, *see* algebras, ~
- testing, 14
- total order, 29
- transitions, 25
 - and statements, 36
 - firing rule, 27
 - guards, 25
 - instances, 27
 - analysis, *see* finding enabled
 - enabled, 27
 - finding enabled, 50–56
 - undefined, 38
 - locality, 30
 - modes, *see* instances
 - nondeterministic, 9
 - reducing the number of, 8
 - valuations, *see* instances
- valuations, *see* transitions, ~
- variables, 21
 - assignments, 22
 - output, 27
 - undefined, 22