

Technical Report

SDL Profiles - Definition and Formal Extraction

R. Grammes R. Gotzhein

November 20, 2006

University of Kaiserslautern

Contents

1	Introduction	4
2	Language Definition of SDL	5
2.1	Specification and Description Language (SDL)	5
2.2	Abstract State Machines	5
2.3	Formal Definition of SDL-2000	6
3	Definition and Extraction of SDL Profiles	8
3.1	SDL Profiles	8
3.2	Outline of the Extraction Approach for SDL Profiles	9
4	Formalisation	11
4.1	Reduction Profile	11
4.2	Formalisation Signature	12
4.3	Formal Reduction of ASMs	13
4.3.1	Formal Reduction of ASM Definitions	13
4.3.2	Macros, Functions and Parameters	14
4.3.3	Formal Reduction of ASM Rules	16
4.3.4	Formal Reduction of ASM Domains	18
4.3.5	Formal Reduction of ASM Expressions	21
5	Consistency of SDL Profiles	30
6	SDL-Profile Tool	32
6.1	Tool Chain	32
6.2	Application of the SDL-Profile Tool	34
7	Related Work	35
8	Conclusions and Outlook	36

Over a period of 30 years, ITU-T's Specification and Description Language (SDL) has matured to a sophisticated formal modelling language for distributed systems and communication protocols. The language definition of SDL-2000, the latest version of SDL, is complex and difficult to maintain. Full tool support for SDL is costly to implement. Therefore, only subsets of SDL are currently supported by tools. These SDL subsets - called *SDL profiles* - already cover a wide range of systems, and are often sufficient in practice. In this report, we present our approach for extracting the formal semantics for SDL profiles from the complete SDL semantics. We then formalise the approach, present our SDL-profile tool, and report on our experiences.

1 Introduction

Over a period of 30 years, ITU-T's Specification and Description Language (SDL) [6, 7, 8] has matured from a simple, informal graphical notation for describing a set of communicating finite state machines to a sophisticated formal modelling technique with graphical syntax, data types, structuring mechanisms, object-oriented features, formal semantics, support for reuse, companion notations, and commercial tool environments. This development has led to an expressive and sophisticated language for a wide range of domains. On the other hand, the language definition of SDL-2000, the latest version of SDL, is complex and difficult to maintain. Full tool support for SDL is costly to implement. Therefore, all commercial tool providers have decided to support subsets of SDL only. These SDL subsets - also called *SDL profiles* in this report - are targeted towards specific domains and companies, where, due to their reduced complexity, they are preferred by engineers.

While the use of SDL profiles is today's state-of-the-practice, their definition is not reflected in the SDL standard. One could argue that this is of no particular importance, since the full language definition covers all possible subsets. However, a drawback is that engineers working with a well-defined SDL profile only are still confronted with the entire language definition. Also, the task of tool builders to show conformance to the language definition is highly complex, in particular if the optimisation potential of a particular SDL profile is to be exploited.

To solve these problems, one could think of defining a separate standard for each SDL profile of interest. This, however, creates other problems arising from the extra work to define and maintain these standards, and of keeping them consistent. In this report, we address these problems and present a formalised, tool-based approach for extracting, for a given SDL profile, the formal semantics from the standardised SDL semantics. In Section 2, we survey the language definition of SDL, in particular, the formal semantics. Section 3 outlines our approach to the extraction of the formal semantics for SDL profiles. In Section 4, the complete formalisation of the approach is presented. Consistency of SDL profiles is visited in Section 5. We present our tool chain for the extraction approach in Section 6, survey related work in Section 7, and draw conclusions in Section 8.

2 Language Definition of SDL

2.1 Specification and Description Language (SDL)

The Specification and Description Language (SDL) [6] is a formal language standardised by the International Telecommunications Union (ITU). It is widely used both in industry and academia. SDL is based on the concept of asynchronously communicating finite state machines, running concurrently or in parallel. SDL provides language constructs for the specification of nested *system structure*, *communication* using channels, signals and signal queues, *behaviour* using extended finite state machines, and *data*.

In 1988, the semantics of SDL was formally defined, upgrading the language to a formal description technique. In 1999, a new version of the language, referred to as SDL-2000, was introduced. Since the formal definition of the semantics was assessed as being too difficult to extend and maintain, a new formal semantics, based on Abstract State Machines, was defined from scratch [2, 9, 10].

2.2 Abstract State Machines

Abstract State Machines (ASMs) [4, 5] are a general model of computation introduced by Yuri Gurevich. They combine declarative concepts of first-order logic with the abstract operational view of distributed transition systems. ASMs are based on many-sorted first-order structures, called *states*. A state consists of a signature containing domain names, function names and relation names, together with an interpretation of these names. A state can be viewed as a memory snapshot of the ASM, where locations - identified by functions and parameter values - are mapped to result values.

ASM-Listing 1 shows the signature for the specification of a graph, consisting of domain names *Node* and *Edge* (defined as a pair of nodes), a 0-ary function name *start* and an unary function names *weight* and *visited*. A state of the ASM defines an interpretation of the names in the signature. *Node* and *Edge* are interpreted as sets of elements, *start* is interpreted as an element of the set *Node* (or the special element **undefined**), *weight* and *visited* are interpreted as functions from elements of set *Edge* to the set of natural numbers and the set of boolean values, respectively.

```
1 domain Node
2 domain Edge =def Node × Node
3
4 start: → Node
5 weight: Edge → ℕ
```

ASM-Listing 1: Vocabulary of a Graph Specification

The computation model of distributed ASMs is based on a set of autonomously operating ASM agents. Starting from an initial state, the agents perform concurrent computations and interact through shared locations of the state. The behaviour of ASM agents is determined by ASM programs, consisting of ASM rules. Complex ASM rules are defined as compositions of guarded update instructions using a small set of rule constructors. From these rules, update sets, i.e. sets of memory locations and new values, are computed. These update sets define state transitions that result from applying all updates simultaneously.

A detailed introduction to Abstract State Machines and the syntax used in this report is given in [2].

2.3 Formal Definition of SDL-2000

In November 2000, the formal semantics of SDL-2000 was officially approved to become part of the SDL language definition. It covers all static and dynamic language aspects, and consists of two major parts (for a detailed survey, see [2]):

- The *static semantics* of SDL defines well-formedness conditions on the concrete syntax of SDL. Transformations map extended features of SDL to core features of the language, reducing the complexity of the dynamic semantics.
- The *dynamic semantics* of SDL defines the dynamic behaviour of well-formed SDL specifications, based on ASMs. At the core of the dynamic semantics is the SDL Virtual Machine (SVM), providing a signal flow model and several types of agents. Agents go through an initialisation phase, creating the nested structure of an SDL system, and an execution phase, forwarding signals and executing state machines. Behaviour primitives form the instruction set of the SVM, defining basic actions such as sending signals, setting timers or calling procedures. A *compilation function* maps actions from transitions in an SDL specification to instructions of the SVM.

ASM Listing 2 shows an excerpt from the SVM, which provides a formal definition for the setting of timers, encapsulated in the ASM rule macro SETTIMER. When evaluated, the macro defines an update of the input queue of an SDL agent, where a possibly existing timer instance is removed, and a new timer instance is appended.

```

1 SETTIMER(tm: TIMER, vSeq: VALUE*, t: [TIME]) ≡
2   let tmi = mk-TimerInst(Self.self, tm, vSeq) in
3     if t = undefined then
4       Self.inport.schedule := insert(tmi, now + tm.duration, delete(tmi, Self.inport.schedule)
5         )
6       tmi.arrival := now + tm.duration
7     else
8       Self.inport.schedule := insert(tmi, t, delete(tmi, Self.inport.schedule))

```

```
8         tmi.arrival := t
9     endif
10 endlet
```

ASM-Listing 2: Setting SDL Timers

The let-rule in line 2 defines a shortcut *tmi* for a new `TIMERINST` tuple, consisting of the process identifier of the executing agent *Self.self*, the type of the timer *tm*, and a sequence of parameters *vSeq*. In case no time *t* is passed as a parameter to the rule macro, *tmi* is inserted into the schedule of the process (line 4) at the current time *now* plus the default duration of the timer type. In the same step, all timer tuples that are structurally equivalent to *tmi* (same type, parameters and agent) are deleted from the schedule. In case a time *t* is passed, the timer instance is inserted into the schedule at *t* (line 7).

3 Definition and Extraction of SDL Profiles

3.1 SDL Profiles

SDL has become a sophisticated and complex language with many language features. This results in a large and extensive language definition. In the formal semantics of SDL-2000, the operational nature of ASMs and the extensive use of modularisation lead to a readable formal semantics definition. However, due to the complexity of the language, the formal semantics is large and requires substantial effort to be understood completely: the dynamic semantics of SDL-2000 consist of more than 3000 lines of ASM specification.

The problem of the complexity of SDL-2000 has been identified, and the definition of simpler sublanguages of SDL has been proposed. One such language is defined by the SDL Task Force as the simplest useful enhanced subset of SDL (referred to as 'SDL+') [17]. This language is implemented by the SAFIRE tool, and here is called SAFIRE. SAFIRE focuses on the state machine aspect of SDL, and enhances it with functionality needed for testing. However, although a formal semantics exists for SDL, none is provided for SAFIRE.

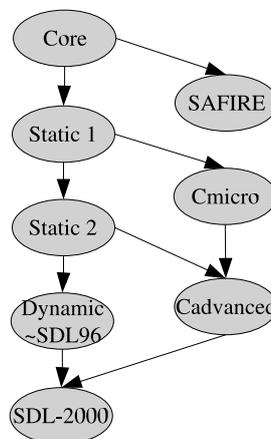


Figure 3.1: Superset Relationship between Language Profiles

SDL profiles are self-contained subsets of SDL, possibly enhanced by further language constructs, targeted towards specific domains and companies. In comparison to the full SDL language definition, SDL profiles have reduced complexity, which supports both engineers and tool builders. A sublanguage like SAFIRE is an SDL profile. Tools for an SDL profile can be developed faster, leading to less expensive tools and enabling code optimisations. Possible SDL profiles could also be derived from the supported features of the code generators Cbasic and Cadvanced in Telelogic Tau.

Apart from being subsets of the complete language, SDL profiles can be subsets of other SDL profiles, forming a hierarchy profiles. We have defined four SDL profiles. The smallest profile is *Core*, which contains a minimal set of features. *Static₁*, *Static₂* and *Dynamic* extend *Core*, each profile adding additional features to the preceding one, *Dynamic* being roughly the equivalent of SDL-96. The subset relationships between different language profiles are shown in Figure 3.1. The SDL profiles *Core*, *Static₁*, *Static₂* and *Dynamic* are supported syntactically by our configurable transpiler ConTraST [1].

3.2 Outline of the Extraction Approach for SDL Profiles

Being language definitions themselves, SDL profiles consist of a formal syntax and a formal semantics. To reduce the overhead of defining SDL profiles and to keep them consistent, a systematic, formalised approach is needed. Basically, there are two ways to define SDL profiles:

- **Composition (bottom-up).** Here, SDL profiles are defined by composing an SDL language core and selected SDL language modules. The *language core* can be understood as the smallest useful SDL subset, for instance, SDL reduced to a set of elementary communicating finite state machines, without any extensions. A *language module* encapsulates a language feature, defining its syntax, semantics, and dependencies to other language modules. Examples of SDL language modules are timers, exceptions, save, and inheritance. Note that language core as well as language modules consist of syntax *and* semantics. For the composition to be feasible, it is crucial that the semantics of modules can be encapsulated and composed with the semantics of the language core and other modules.
- **Extraction (top-down).** Here, SDL profiles are defined by extracting the profile definition from the complete SDL language definition. As in the composition approach, SDL language modules, each consisting of a set of SDL language constructs, can be identified. To obtain a particular SDL profile, these language modules are then removed from the complete language definition. Different from the composition approach, it is not necessary that the semantics of modules can be encapsulated. Instead, it suffices to characterise modules by their language constructs. With this information, it is straightforward to identify corresponding grammar rules, and to reduce or remove them. A major problem to be solved is the reduction of the formal semantics.

From a methodological point of view, the composition approach seems more appealing. However, there is the difficulty of encapsulating the formal semantics of SDL language modules such that composition is supported, which we have not been able to overcome. Therefore, we have investigated the extraction approach, with quick first results. Meanwhile, we have further developed, consolidated, and formalised the extraction approach, and have built a sophisticated tool chain.

To extract the formal semantics of an SDL profile from the complete SDL semantics, we have considered two approaches:

- **ASM rule coverage.** With each SDL profile, an ASM rule coverage comprising all ASM rules of the SVM that may be evaluated in some execution of some SDL specification written in that SDL profile can be associated. While this approach is semantically sound, it is practically infeasible. For a given SDL specification, the concurrent, non-deterministic nature of the SVM may lead to a very large number of possible executions. Furthermore, the number of SDL specifications that can be written in a given SDL profile is extremely large. Therefore, the worst-case complexity of an algorithm for ASM rule coverage is far too high to be of any use for practical purposes.
- **dead ASM rule recognition.** Instead of computing the ASM rule coverage of a set of SDL specifications, we can develop safe criteria to recognise ASM rules that are never evaluated for a given SDL profile. For instance, if the SDL language module *timer* is to be removed, we can safely remove all ASM rules that are used for setting and resetting SDL timers, including the corresponding ASM domains, functions, and relations. It is important here that ASM dead rule recognition works in a conservative way, i.e. ASM rules must only be removed if it can be formally proven that they are not evaluated for a given SDL profile. The degree of reduction that can be achieved this way thus depends on the completeness of the criteria that can be defined. Unlike the ASM rule coverage approach, dead code recognition is practically feasible. Therefore, we have followed this approach, and will present safe criteria as well as some heuristics below.

4 Formalisation

4.1 Reduction Profile

SDL profiles characterise subsets of the set of valid SDL specifications, by defining subsets of the concrete and abstract syntax of SDL. The abstract syntax of SDL influences the dynamic semantics, which is the focus of our work, in two ways:

- The abstract syntax yields part of the SVM data structure (ASM signature, see Figure 4.1). For each element of the abstract grammar, a domain of the same name is introduced in the ASM signature. E.g., the following non-terminals of the abstract grammar, which are only relevant for SDL specifications with timers, are also domains in the signature of the ASM: *Timer-name*, *Timer-identifier*, *Timer-definition*, *Timer-active-expression*, *Set-node*, and *Reset-node*.
- In the case of SDL actions (e.g., assignments, setting timers), a compilation function maps parts of the abstract syntax to domains of the formal semantics definition that form the SVM. E.g., the compilation of a *Set-node* in the abstract syntax tree leads to the creation of an element of the domain SET in the ASM signature.

Following the extraction approach, we remove SDL language modules from the formal language definition. Language modules consist of sets of SDL language constructs, and their corresponding grammar rules. These grammar rules are removed from the formal syntax definition. Furthermore, they form the starting point for the reduction of the formal semantics definition (see Figure 4.1). Starting from the removed parts of the formal syntax definition, we can identify corresponding domains in the ASM signature, as described above. These domains are empty in the initial state of the SVM, and, since they are not modified by the SVM, will be empty in all reachable states, too. This observation is fundamental for recognising dead ASM rules of the SVM.

Apart from domains corresponding to elements of the abstract grammar of a language module, other domains, functions and predicates in the SVM signature correspond to specific language modules. E.g., *SignalSaved* is a predicate that corresponds to the *save* feature in SDL. If it holds, the signal being examined is not discarded, if no valid transition is found. These elements of the SVM signature are removed in addition to domains corresponding to elements of the abstract grammar. However, we need to prove that these elements are not needed for the given SDL profile. E.g., we can show that *SignalSaved* is always false if *save-signalset* is empty.

In order to perform ASM dead code recognition, we specify all parts of the ASM signature that correspond to language modules not included in the SDL profile in a

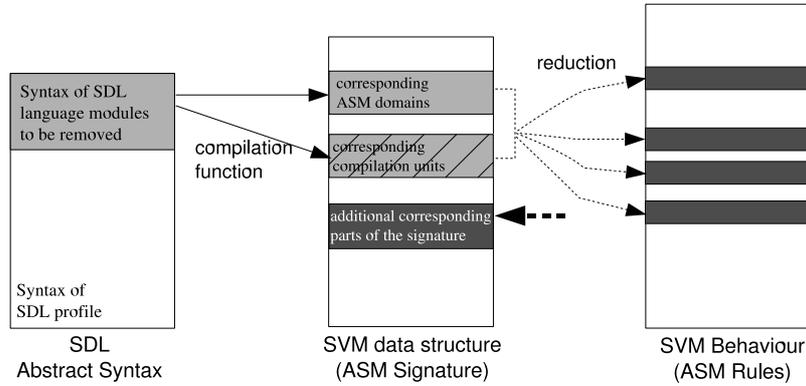


Figure 4.1: Concept of the Extraction Process

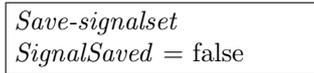


Figure 4.2: Reduction Profile for 'Save' Feature

reduction profile. The reduction profile is a list of domains, functions and predicates from the SVM signature to be removed in the extraction process. Additionally, we specify a default value (**true** or **false**) for predicates. These elements are removed from the formal language definition according to the rules formally defined in the following sections.

Figure 4.2 shows the smallest possible reduction profile corresponding to a language feature. It specifies all grammar elements and predicates used to defer the consumption of input signals.

4.2 Formalisation Signature

We now formalise our approach for extracting the formal semantics of SDL profiles from the complete SDL semantics. The formalisation gives a precise definition of the removal process, which leads to deterministic results, and provides the foundation for tool support for the removal process. Finally, a formal definition is necessary in order to make precise statements about the consistency of SDL profiles. Since the formal syntax definition can be easily defined in a modular fashion, making its reduction

straightforward, we focus on the reduction of the formal semantics definition.

For the formal definition of the extraction process, we have decided to use a functional approach, defining functions that recursively map the original formal semantics to the reduced formal semantics. These functions are based on a concrete grammar for Abstract State Machines [3]. The input of the reduction is the SDL formal semantics definition from [10] and a reduction profile r , as described in Section 4.1.

To formalise the extraction, we define a function *remove*, which maps a term from the grammar G of ASMs and a set of variables V - an initially empty set of locally undefined variables from the ASM formal semantics - to a reduced term from the grammar G . Additionally, we introduce three *mutually exclusive* binary predicates, namely *undefined_r*, *true_r* and *false_r*, that control the reduction. Predicates *undefined_r*, *true_r* and *false_r* form the interface to the reduction profile. Furthermore, *true_r* and *false_r* define heuristics to determine if an expression is **true** (**false**) in every reachable state.

$$\begin{aligned} \text{remove}_r &: G \times V \rightarrow G \\ \text{undefined}_r &: G \times V \rightarrow \text{Boolean} \\ \text{true}_r &: G \times V \rightarrow \text{Boolean} \\ \text{false}_r &: G \times V \rightarrow \text{Boolean} \end{aligned}$$

The *remove* function is defined on all elements of the grammar G . Predicates *true_r* and *false_r* are defined on boolean and first-order logic expressions, and predicate *undefined_r* on all expressions. In the following sections, we omit the index r from the predicates.

The function *remove* is defined recursively - a given term is mapped to a new term by applying the mapping defined by *remove* to the subterms. In case none of the predicates *undefined*, *true* and *false* holds, the current term is not reduced any further. This assures in particular that *remove* corresponds to the identical mapping if the reduction profile r is empty. In other cases, subterms can be replaced or omitted depending on which of the predicates hold.

4.3 Formal Reduction of ASMs

4.3.1 Formal Reduction of ASM Definitions

This section describes the *remove* function for ASM definitions, namely domain, function and rule macro definitions. The *remove* function does not remove a definition completely. If, by the reduction process, the definition becomes trivial or is not referenced anymore, it can be removed in a subsequent cleanup step. The following variables are used to define the removal:

$D, D_1, D_2 \in \text{domain}$, $\text{dn} \in \text{DomainName}$, $\text{exp} \in \text{formula}$, $R, R_1, R_2 \in \text{rule}$, $\text{ps} \in \text{paramSeq}$

Domain definitions are not affected by removal, unless they are derived definitions. For derived domain definitions, removal continues with the domain expression that defines the derived domain.

$$\begin{aligned} \text{remove}(\text{mode } \mathbf{domain} \text{ dn}, \mathcal{V}) &= \text{mode } \mathbf{domain} \text{ dn} \\ \text{remove}(\mathbf{domain} \text{ dn}, \mathcal{V}) &= \mathbf{domain} \text{ dn} \\ \text{remove}(\text{dn} =_{\text{def}} D, \mathcal{V}) &= \text{dn} =_{\text{def}} \text{remove}(D, \mathcal{V}) \end{aligned}$$

For function definitions, removal continues with the function signature to remove any undefined domains. For derived functions, removal also continues with the formula that defines the function, taking into account all formal parameters that were removed.

$$\begin{aligned} \text{remove}(\text{mode } f \text{ ':' } D_1 \rightarrow D_2, \mathcal{V}) &= \text{mode } f \text{ ':' } \text{remove}(D_1, \mathcal{V}) \rightarrow \text{remove}(D_2, \mathcal{V}) \\ \text{remove}(\text{mode } f \text{ ':' } \rightarrow D_2, \mathcal{V}) &= \text{mode } f \text{ ':' } \rightarrow \text{remove}(D_2, \mathcal{V}) \\ \text{remove}(f \text{ ':' } D_1 \rightarrow D_2, \mathcal{V}) &= f \text{ ':' } \text{remove}(D_1, \mathcal{V}) \rightarrow \text{remove}(D_2, \mathcal{V}) \\ \text{remove}(f \text{ ':' } \rightarrow D_2, \mathcal{V}) &= f \text{ ':' } \rightarrow \text{remove}(D_2, \mathcal{V}) \end{aligned}$$

$$\begin{aligned} \text{remove}(f \text{ ':' } D =_{\text{def}} \text{exp}, \mathcal{V}) &= f \text{ ':' } \text{remove}(D, \mathcal{V}) =_{\text{def}} \text{remove}(\text{exp}, \mathcal{V}) \\ \text{remove}(f(\text{ps}) \text{ ':' } D =_{\text{def}} \text{exp}, \mathcal{V}) &= \\ f(\text{remove}(\text{ps}, \mathcal{V})) \text{ ':' } \text{remove}(D, \mathcal{V}) &=_{\text{def}} \text{remove}(\text{exp}, \mathcal{V} \cup \text{remfpar}(\text{ps})) \end{aligned}$$

A macro definition consists of a rule name, a sequence of formal parameters and a rule body. Domains of formal parameters may be undefined, and the corresponding parameters must be removed. Removal continues with the rule body and the undefined formal parameters added to the list of undefined variables.

$$\begin{aligned} \text{remove}(\mathbf{RuleName} \equiv R, \mathcal{V}) &= \mathbf{RuleName} \equiv \text{remove}(R, \mathcal{V}) \\ \text{remove}(\mathbf{RuleName}(\text{ps}) \equiv R, \mathcal{V}) &= \\ \mathbf{RuleName}(\text{remove}(\text{ps}, \mathcal{V})) \equiv \text{remove}(R, \mathcal{V} \cup \text{remfpar}(\text{ps})) & \end{aligned}$$

Removal on constraints equates to removal on the constraint formula. Removal of program definitions continues with removal of the rule body of the program.

$$\begin{aligned} \text{remove}(\mathbf{constraint} \text{ exp}, \mathcal{V}) &= \mathbf{constraint} \text{ remove}(\text{exp}, \mathcal{V}) \\ \text{remove}(\mathbf{initially} \text{ exp}, \mathcal{V}) &= \mathbf{initially} \text{ remove}(\text{exp}, \mathcal{V}) \end{aligned}$$

$$\begin{aligned} \text{remove}(\mathbf{ProgramName} \text{ ':' } R, \mathcal{V}) &= \mathbf{ProgramName} \text{ ':' } \text{remove}(R, \mathcal{V}) \\ \text{remove}(\mathbf{ProgramName} \text{ ':' }, \mathcal{V}) &= \mathbf{ProgramName} \text{ ':' } \end{aligned}$$

4.3.2 Macros, Functions and Parameters

This section describes the removal of formal parameters of rule macros and functions, and the removal of corresponding parameters from calls to these macros and functions. The following variables are used in this section, in addition to the ones introduced above:

$fcs \in \text{formulaCommaSeq}, n, p \in \mathbb{N}$

Formal parameters are removed from a list of formal parameters if *undefined* holds for their domain (the type). Removal of formal parameters starts with the rightmost parameter.

$$\begin{aligned} \text{remove}(\text{ps } ', ' x ': ' D, \mathcal{V}) = & \\ & \begin{array}{ll} \text{remove}(\text{ps}, \mathcal{V}) & \text{iff } \text{undefined}(D) \\ \text{remove}(\text{ps}, \mathcal{V}) ', ' x ': ' \text{remove}(D, \mathcal{V}) & \text{else} \end{array} \end{aligned}$$

$$\begin{aligned} \text{remove}(x ': ' D, \mathcal{V}) = & \\ \text{''} & \text{iff } \text{undefined}(D) \\ x ': ' \text{remove}(D, \mathcal{V}) & \text{else} \end{aligned}$$

The function *numfpar* counts the number of formal parameters in a formal parameter sequence. The function is used when removing parameters from a parameter sequence (see below).

$$\begin{aligned} \text{numfpar}(\text{fcs } ', ' \text{exp}) &= \text{numfpar}(\text{fcs}) + 1 \\ \text{numfpar}(\text{exp}) &= 1 \\ \text{numfpar}(\text{ps } ', ' x ': ' D) &= \text{numfpar}(\text{ps}) + 1 \\ \text{numfpar}(x ': ' D) &= 1 \end{aligned}$$

The function *remfpar* returns a set of names of formal parameters. The set includes all names of a formal parameter sequence for which *undefined* holds for the corresponding domain.

$$\begin{aligned} \text{remfpar}(\text{ps } ', ' x ': ' D) = & \\ \{x\} \cup \text{remfpar}(\text{ps}) & \text{iff } \text{undefined}(D) \\ \text{remfpar}(\text{ps}) & \text{else} \end{aligned}$$

$$\begin{aligned} \text{remfpar}(x ': ' D) = & \\ \{x\} & \text{iff } \text{undefined}(D) \\ \{\} & \text{else} \end{aligned}$$

Formal parameters removed in a macro definition must be removed from the argument lists of macro calls. *count* assigns a code to a macro that describes which parameters have been removed. The code function gets a list of parameters and a number n (initially the number of arguments minus one) as arguments. If the domain of the rightmost argument is undefined, 2^n is added to the code of the remaining parameters with the number $n - 1$. E.g. for a sequence of four parameters, with the first and the third undefined, the code is $2^2 + 2^0 = 5$.

$$\begin{aligned} \text{code}(\text{ps } ', ' x ': ' D, n) = & \\ \text{code}(\text{ps}, n - 1) + 2^n & \text{iff } \text{undefined}(D) \\ \text{code}(\text{ps}, n - 1) & \text{else} \end{aligned}$$

$$\begin{aligned} \text{code}(x \text{ ':' } D, n) = & \\ & \begin{array}{ll} 2^n & \text{iff } \text{undefined}(D) \\ 0 & \text{else} \end{array} \end{aligned}$$

$$\text{count}(\text{MacroName}) = \text{code}(\text{ps}, \text{numfpar}(\text{ps}) - 1)$$

The function *removepar* removes arguments from a macro call corresponding to undefined formal parameters in the macro definition. The number n corresponds to the position of the argument (initially the number of arguments minus one), the number p to the code of the macro definition. If p is larger than 2^n , the argument is removed and removal is continued with the remaining parameters.

$$\begin{aligned} \text{removepar}(\text{fcs } \text{'}, \text{' exp}, n, p) = & \\ & \begin{array}{ll} \text{removepar}(\text{fcs}, n - 1, p - 2^n) & \text{iff } p - 2^n > 0 \\ \text{removepar}(\text{fcs}, n - 1, p) \text{ ', ' exp} & \text{else} \end{array} \end{aligned}$$

$$\begin{aligned} \text{removepar}(\text{exp}, n, p) = & \\ \text{"} & \text{iff } p = 2^n \text{ // } n \text{ should be } 0 \\ \text{exp} & \text{else} \end{aligned}$$

A sequence of formulas is undefined if each formula in the sequence is undefined.

$$\text{undefined}(\text{fcs}, \text{exp}, \mathcal{V}) \text{ iff } \text{undefined}(\text{fcs}, \mathcal{V}) \vee \text{undefined}(\text{exp}, \mathcal{V})$$

4.3.3 Formal Reduction of ASM Rules

Rules specify transitions between states of the ASM. The basic rule is the *update rule*, which updates a location of the state to a new value. All together, there are seven kinds of rules for ASMs, for all of which we have formalised the reduction.

The left hand side of an update rule specifies a location of the ASM. The location consists of a function f from the ASM signature and a tuple of elements fcs . If *undefined* holds for either, the location lies outside the scope of the reduced ASM, and the update rule is omitted. If *undefined* holds for the expression on the right hand side of the update rule, we remove the update rule, retaining the previous value of the location.

$$\begin{aligned} \text{remove}(f(\text{fcs}) := \text{exp}, \mathcal{V}) = & \\ \text{skip} & \text{iff } \text{undefined}(f, \mathcal{V}) \vee \text{undefined}(\text{fcs}, \mathcal{V}) \vee \\ & \text{undefined}(\text{exp}, \mathcal{V}) \\ f(\text{remove}(\text{fcs}, \mathcal{V})) := \text{remove}(\text{exp}, \mathcal{V}) & \text{else} \end{aligned}$$

The mapping of the **if**-rule (see below) depends on which predicate holds for the guard *exp* of the rule. If the guard always evaluates to **true** (**false**), the **if**-rule can be

omitted, and removal continues with subrule R_1 (R_2). If the guard is undefined, the rule is syntactically incorrect, and should not be reachable¹. If none of the predicates hold, the removal is applied recursively to the guard and the subrules of the **if**-rule, leaving the rule itself intact.

$$\begin{aligned}
\text{remove}(\mathbf{if\ exp\ then\ } R_1 \mathbf{\ else\ } R_2 \mathbf{\ endif,\ } \mathcal{V}) = & \\
\text{remove}(R_1, \mathcal{V}) & \quad \text{iff } \text{true}(\text{exp}, \mathcal{V}) \\
\text{remove}(R_2, \mathcal{V}) & \quad \text{iff } \text{false}(\text{exp}, \mathcal{V}) \\
\text{skip} & \quad \text{iff } \text{undefined}(\text{exp}, \mathcal{V}) \\
\mathbf{if\ } \text{remove}(\text{exp}, \mathcal{V}) \mathbf{\ then\ } \text{remove}(R_1, \mathcal{V}) & \quad \text{else} \\
\mathbf{\ else\ } \text{remove}(R_2, \mathcal{V}) \mathbf{\ endif} &
\end{aligned}$$

The **let**-rule is a shortcut that binds the evaluation result of an expression in the current state to a variable, which can be used inside the **let**-rule. In case the expression exp is undefined, so is the variable x . The result is the mapping of the contained rule R , with the variable x included in the set of locally undefined names \mathcal{V} . The result of the removal is the same as if the expression exp had been used directly in the rule R instead of the variable x .

$$\begin{aligned}
\text{remove}(\mathbf{let\ } x = \text{exp} \mathbf{\ in\ } R \mathbf{\ endlet,\ } \mathcal{V}) = & \\
\text{remove}(R, \mathcal{V} \cup \{x\}) & \quad \text{iff } \text{undefined}(\text{exp}, \mathcal{V}) \\
\mathbf{let\ } x = \text{remove}(\text{exp}, \mathcal{V}) \mathbf{\ in\ } \text{remove}(R, \mathcal{V}) & \quad \text{else} \\
\mathbf{endlet} &
\end{aligned}$$

In **let**-rules, we can additionally specify a type for the variable x in form of a domain D . In that case, we can remove the **let**-rule in the same manner as if the expression exp were undefined, if the domain D is undefined.

$$\begin{aligned}
\text{remove}(\mathbf{let\ } x : D = \text{exp} \mathbf{\ in\ } R \mathbf{\ endlet,\ } \mathcal{V}) = & \\
\text{remove}(R, \mathcal{V} \cup \{x\}) & \quad \text{iff } \text{undefined}(\text{exp}, \mathcal{V}) \vee \text{undefined}(D) \\
\mathbf{let\ } x : \text{remove}(D, \mathcal{V}) = \text{remove}(\text{exp}, \mathcal{V}) \mathbf{\ in\ } & \quad \text{else} \\
\text{remove}(R, \mathcal{V}) \mathbf{endlet} &
\end{aligned}$$

The **extend**-rule dynamically imports a fresh ASM element from the reserve (an infinite store of unused ASM elements), binding it to a variable x in the context of the subrule R and including it in the ASM domain dn (given by name). In case the domain name dn is undefined, i.e. has been removed from the ASM signature, the **extend**-rule can be omitted, since elements of domain dn belong to a removed feature. However, the subrule R might still contain parts not related to this feature - although it would be a better style to move these parts outside the **extend**-rule. Therefore, the subrule is not omitted by default, but replaced with its mapping by the remove function, including the now unbound variable x in the set of locally undefined variables. This leads to all occurrences of x being removed from the rule R .

¹This is a proof obligation that we have to verify manually. However, so far this has only occurred in very few cases, which were the result of errors in the reduction profile.

$$\begin{aligned}
\text{remove}(\text{extend dn with } x R \text{ endextend}, \mathcal{V}) = & \\
\text{remove}(R, \mathcal{V} \cup \{x\}) & \quad \text{iff} \quad \text{undefined}(\text{dn}, \mathcal{V}) \\
\text{extend dn with } x \text{ remove}(R, \mathcal{V}) \text{ endextend} & \quad \text{else}
\end{aligned}$$

The **choose**-rule nondeterministically takes an element from the finite set defined by the constraint *exp* and binds it to the variable *x*. If no element satisfies the constraint, as in the case where *false* holds, choose is equivalent to skip. Furthermore, if *undefined* holds for the constraint, we assume that no element matches it. If *true* holds for the constraint, the **choose**-rule is invalid since it ranges over a potentially infinite set.

$$\begin{aligned}
\text{remove}(\text{choose } x : \text{exp } R \text{ endchoose}, \mathcal{V}) = & \\
\text{skip} & \quad \text{iff} \quad \text{false}(\text{exp}, \mathcal{V}) \vee \text{true}(\text{exp}, \mathcal{V}) \vee \\
& \quad \text{undefined}(\text{exp}, \mathcal{V}) \\
\text{choose } x : \text{remove}(\text{exp}, \mathcal{V}) \text{ remove}(R, \mathcal{V}) & \quad \text{else} \\
\text{endchoose} &
\end{aligned}$$

Using typed ASMs, it is sensible to restrict the element *x* to a domain *D* as the type of *x*. In the formal semantics of SDL-2000, all constraints of **choose**-rules have the form "*x* ∈ *D* ∧ *constraint*" (*constraint* being optional). From the definition of *true* and *undefined* on expressions (see 4.3.5), we can conclude that these predicates will not hold for a constraint of this form.

The **do forall**-rule performs a parallel update of the state, firing the rule *R* with *x* bound to the element *a*, for all *a* ∈ {*x* | *exp*}, {*x* | *exp*} being a finite set. Removal follows the same principles as with **choose**, as both **choose**-rule and **do forall**-rule use elements from a finite set defined by a constraint. If *false* holds for the constraint, the rule is equivalent to skip. Predicates *true* and *undefined* do not hold if *exp* has the form "*x* ∈ *D* ∧ *constraint*".

$$\begin{aligned}
\text{remove}(\text{do forall } x \text{ :? exp } R \text{ enddo}, \mathcal{V}) = & \\
\text{skip} & \quad \text{iff} \quad \text{false}(\text{exp}, \mathcal{V}) \vee \text{true}(\text{exp}, \mathcal{V}) \vee \\
& \quad \text{undefined}(\text{exp}, \mathcal{V}) \\
\text{do forall } x \text{ :? remove}(\text{exp}, \mathcal{V}) \text{ remove}(R, \mathcal{V}) & \quad \text{else} \\
\text{enddo} &
\end{aligned}$$

Rule blocks in ASMs are fired in parallel. A sequence of rule blocks is broken down to the mappings of the sub-rule blocks. This may result in a sequence of skip-rules which can be reduced to a single skip. However, this is not part of the remove mapping, but can be done in a subsequent optimisation step.

$$\text{remove}(R_1 R, \mathcal{V}) = \text{remove}(R_1, \mathcal{V}) \text{ remove}(R, \mathcal{V})$$

4.3.4 Formal Reduction of ASM Domains

This section defines the *remove* function for expressions describing ASM domains, e.g. union or tuple domains. The following variables are used, in addition to the variables defined in previous sections:

$s, s_1, s_2 \in \text{simpledomain}, t \in \text{tupledomain}, u \in \text{uniondomain}, \text{ics} \in \text{itemCommaSeq}.$

A domain is removed if it is undefined. A union or product of several domains is removed if all of the domains are undefined, otherwise only the undefined subdomains are removed.

$remove(dn, \mathcal{V}) =$		
nodomain	iff	$undefined(dn, \mathcal{V})$
dn		else
$remove(dn \text{ mod}, \mathcal{V}) =$		
nodomain	iff	$undefined(dn)$
$dn \text{ mod}$		else
$remove(\underline{(D)}, \mathcal{V}) =$		
nodomain	iff	$undefined(D, \mathcal{V})$
$(remove(D, \mathcal{V}))$		else
$remove(\underline{(D)}\text{-set}, \mathcal{V}) =$		
nodomain	iff	$undefined(D, \mathcal{V})$
$(remove(D, \mathcal{V}))\text{-set}$		else
$remove(\underline{[D]}, \mathcal{V}) =$		
nodomain	iff	$undefined(D, \mathcal{V})$
$[remove(D, \mathcal{V})]$		else
$remove(\underline{[D]} \text{ mod}, \mathcal{V}) =$		
nodomain	iff	$undefined(D, \mathcal{V})$
$[remove(D, \mathcal{V})] \text{ mod}$		else
$remove(D_1 \rightarrow D_2, \mathcal{V}) =$		
nodomain	iff	$undefined(D_1, \mathcal{V}) \vee undefined(D_2, \mathcal{V})$
$remove(D_1, \mathcal{V}) \rightarrow remove(D_2, \mathcal{V})$		else
$remove(s_1 \times s_2, \mathcal{V}) =$		
"	iff	$undefined(s_1 \times s_2, \mathcal{V})$
$remove(s_1, \mathcal{V})$	iff	$undefined(s_2, \mathcal{V})$
$remove(s_2, \mathcal{V})$	iff	$undefined(s_1, \mathcal{V})$
$remove(s_1, \mathcal{V}) \times remove(s_2, \mathcal{V})$		else
$remove(t \times s, \mathcal{V}) =$		
"	iff	$undefined(t \times s, \mathcal{V})$
$remove(t, \mathcal{V})$	iff	$undefined(s, \mathcal{V})$
$remove(s, \mathcal{V})$	iff	$undefined(t, \mathcal{V})$
$remove(t, \mathcal{V}) \times remove(s, \mathcal{V})$		else
$remove('(' \ ')', \mathcal{V}) = '(\)'$		
$remove(s_1 \cup s_2, \mathcal{V}) =$		
"	iff	$undefined(s_1 \cup s_2, \mathcal{V})$
$remove(s_1, \mathcal{V})$	iff	$undefined(s_2, \mathcal{V})$
$remove(s_2, \mathcal{V})$	iff	$undefined(s_1, \mathcal{V})$
$remove(s_1, \mathcal{V}) \cup remove(s_2, \mathcal{V})$		else

$$\begin{aligned}
\text{remove}(u \cup s, \mathcal{V}) = & \\
\text{”} & \text{iff } \text{undefined}(u \cup s, \mathcal{V}) \\
\text{remove}(u, \mathcal{V}) & \text{iff } \text{undefined}(s, \mathcal{V}) \\
\text{remove}(s, \mathcal{V}) & \text{iff } \text{undefined}(u, \mathcal{V}) \\
\text{remove}(u, \mathcal{V}) \cup \text{remove}(s, \mathcal{V}) & \text{else} \\
\text{remove}(\text{'\{ ics \}'}, \mathcal{V}) = & \text{if } \text{undefined}(\text{ics}, \mathcal{V}) \text{ then ” else ’\{ remove(ics, } \mathcal{V}) \}’ \text{ endif}
\end{aligned}$$

The predicate *undefined* on domains specifies if a domain expression is undefined, given the basic domains that have been removed. A domain expression is undefined if the domain name it contains is undefined. In case of union and product of two domains, both domains must be undefined - if only one is undefined, a valid domain definition can be extracted by removing the undefined domain from the expression.

$$\begin{aligned}
\text{undefined}(\text{dn}_*, \mathcal{V}) & \text{iff } \text{undefined}(\text{dn}) \\
\text{undefined}(\text{dn}_{**}, \mathcal{V}) & \text{iff } \text{undefined}(\text{dn}) \\
\text{undefined}(\text{dn}_+, \mathcal{V}) & \text{iff } \text{undefined}(\text{dn}) \\
\text{undefined}(\text{dn-set}, \mathcal{V}) & \text{iff } \text{undefined}(\text{dn}) \\
\text{undefined}(\underline{(D)}, \mathcal{V}) & \text{iff } \text{undefined}(D, \mathcal{V}) \\
\text{undefined}(\underline{(D)}\text{-set}, \mathcal{V}) & \text{iff } \text{undefined}(D, \mathcal{V}) \\
\text{undefined}(\underline{[D]}, \mathcal{V}) & \text{iff } \text{undefined}(D, \mathcal{V}) \\
\text{undefined}(\underline{[D]}_*, \mathcal{V}) & \text{iff } \text{undefined}(D, \mathcal{V}) \\
\text{undefined}(\underline{[D]}_+, \mathcal{V}) & \text{iff } \text{undefined}(D, \mathcal{V}) \\
\text{undefined}(D_1 \rightarrow D_2, \mathcal{V}) & \text{iff } \text{undefined}(D_1, \mathcal{V}) \vee \text{undefined}(D_2, \mathcal{V})
\end{aligned}$$

$$\begin{aligned}
\text{undefined}(s_1 \times s_2, \mathcal{V}) & \text{iff } \text{undefined}(s_1, \mathcal{V}) \wedge \text{undefined}(s_2, \mathcal{V}) \\
\text{undefined}(t \times s, \mathcal{V}) & \text{iff } \text{undefined}(t, \mathcal{V}) \wedge \text{undefined}(s, \mathcal{V}) \\
\text{undefined}(\text{'('}, \mathcal{V}) & \text{iff } \text{false}
\end{aligned}$$

$$\begin{aligned}
\text{undefined}(s_1 \cup s_2, \mathcal{V}) & \text{iff } \text{undefined}(s_1, \mathcal{V}) \wedge \text{undefined}(s_2, \mathcal{V}) \\
\text{undefined}(u \cup s, \mathcal{V}) & \text{iff } \text{undefined}(u, \mathcal{V}) \wedge \text{undefined}(s, \mathcal{V})
\end{aligned}$$

$$\begin{aligned}
\text{undefined}(\text{'\{ics\}'}, \mathcal{V}) & \text{iff } \text{undefined}(\text{ics}, \mathcal{V}) \\
\text{undefined}(\text{ics}', x, \mathcal{V}) & \text{iff } x \in \mathcal{V} \\
\text{undefined}(\text{ics}', \text{kw}, \mathcal{V}) & \text{iff } \text{undefined}(\text{ics}, \mathcal{V}) \wedge \text{undefined}(\text{kw}) \\
\text{undefined}(\text{ics}', \text{kw1 kw2}, \mathcal{V}) & \text{iff } \text{undefined}(\text{ics}, \mathcal{V}) \wedge \text{undefined}(\text{kw1}) \wedge
\end{aligned}$$

undefined(kw2)

undefined(ics', 'Literal, V) iff false

4.3.5 Formal Reduction of ASM Expressions

Expressions are terms over the signature of the SVM. Additionally, ASMs include common mathematical structures like boolean algebra, or natural numbers. Our formal reduction covers all operations defined in [3]. In the truth tables defined in this section, we use the following shortcuts:

T Predicate *true* holds
 F Predicate *false* holds
 U Predicate *undefined* holds
 - $\neg T \wedge \neg F \wedge \neg U$

Furthermore, we use the variables $e, e_1, e_2, e_3 \in \text{formula}$, $\text{nseq} \in \text{nameCommaSeq}$ and $\text{pcs} \in \text{primaryCommaSeq}$.

Boolean Operators

Boolean Operators take boolean expressions as arguments, therefore the predicates *true*, *false* and *undefined* apply. With binary boolean operators, we have to consider sixteen different combinations of predicates holding for the subexpressions - four for each subexpression. In order to improve readability, we combine the definitions of *true*, *false*, *undefined* and *remove* for boolean operators in a four-valued truth table. Valid boolean expressions always evaluate to either **true** or **false**. Therefore, it is undesirable that the predicate *undefined* holds for such an expression. However, this can not be avoided in every case.

We define truth tables for all boolean operators from the concrete syntax of ASMs: negation (\neg), disjunction (\vee), conjunction (\wedge), implication (\rightarrow) and equivalence (\leftrightarrow). In order to ensure consistent results, we derive to definition of conjunction, implication and equivalence from the definitions of negation and disjunction. For the predicates *true* and *false*, the subtables match the truth tables for the corresponding boolean operators with the truth values **true** and **false**, respectively. If all subexpressions of the operator are undefined, so is the composite expression.

The truth table for the negation directly follows from these considerations. In case no predicate holds for the boolean expression e_1 ($-$), removal maps to the original term, with removal applied to the subexpression e_1 .

e_1	\neg	T	F	U	-
		F	T	U	$\neg e_1$

If *true* holds for one of the subexpressions e_1 or e_2 , *true* holds for $e_1 \vee e_2$. If *undefined* holds for one of the subexpressions, it is omitted and the result depends exclusively on the other subexpression. If *false* holds for one of the subexpressions, the subexpression is omitted but can still influence the final result (as in the case *false* and *undefined*).

e_1	e_2				
	\vee	T	F	U	-
T	T	T	T	T	T
F	T	T	F	F	e_1
U	T	T	F	U	e_1
-	T	T	e_2	e_2	$e_1 \vee e_2$

$e_1 \wedge e_2$ is defined as $\neg(\neg e_1 \vee \neg e_2)$, the truth table is derived accordingly.

e_1	e_2				
	\wedge	T	F	U	-
T	T	T	F	T	e_1
F	F	F	F	F	F
U	T	T	F	U	e_1
-	e_2	e_2	F	e_2	$e_1 \wedge e_2$

$e_1 \rightarrow e_2$ is defined as $\neg e_1 \vee e_2$, the truth table is derived accordingly.

e_1	e_2				
	\rightarrow	T	F	U	-
T	T	T	T	T	T
F	F	F	T	F	$\neg e_1$
U	F	F	T	U	$\neg e_1$
-	e_2	e_2	T	$\neg e_1$	$e_1 \rightarrow e_2$

$e_1 \leftrightarrow e_2$ is defined as $(e_1 \rightarrow e_2) \wedge (e_2 \rightarrow e_1)$, the truth table is derived accordingly.

$$\begin{aligned}
\text{true}(\mathbf{if\ } e \mathbf{\ then\ } e_1 \mathbf{\ else\ } e_2 \mathbf{\ endif}) &\text{ iff } (\text{true}(e) \wedge \text{true}(e_1)) \vee \\
&\quad (\text{false}(e) \wedge \text{true}(e_2)) \\
\text{false}(\mathbf{if\ } e \mathbf{\ then\ } e_1 \mathbf{\ else\ } e_2 \mathbf{\ endif}) &\text{ iff } (\text{true}(e) \wedge \text{false}(e_1)) \vee \\
&\quad (\text{false}(e) \wedge \text{false}(e_2)) \\
\text{undefined}(\mathbf{if\ } e \mathbf{\ then\ } e_1 \mathbf{\ else\ } e_2 \mathbf{\ endif}) &\text{ iff } (\text{true}(e) \wedge \text{undefined}(e_1)) \vee \text{undefined}(e) \\
&\quad \vee (\text{false}(e) \wedge \text{undefined}(e_2)) \vee \\
&\quad (\text{undefined}(e_1) \wedge \text{undefined}(e_2))
\end{aligned}$$

Quantification

Quantification consists of two subexpressions - the expression e_1 representing the set of elements in the range of the quantification, and the boolean-valued expression e_2 as the predicate. In the context of quantification, we interpretate e_1 as the empty set if the predicate *undefined* holds.

$$Qx \in e_1 : e_2, Q \in \{\forall, \exists, \exists_1\}$$

Quantification over an empty set (i.e., *undefined* holds for e_1) is always true in case of universal quantification, and always false in case of existential quantification. Furthermore, universal quantification is always true of the boolean expression e_2 is always true, and existential quantification is always false iff e_2 is always false. This leads to the following definitions of the *remove* function and respective predicates.

	e_1				
e_2	\forall	T	F	U	-
	U	T	T	T	T
	-	T	-	U	-

$remove(\forall nseq \in e_1' : 'e_2, \mathcal{V}) =$
true iff $undefined(e_1, \mathcal{V}) \vee true(e_2, \mathcal{V})$
undefined iff $\neg undefined(e_1, \mathcal{V}) \wedge undefined(e_2, \mathcal{V})$
 $\forall nseq \in remove(e_1, \mathcal{V})' : 'remove(e_2, \mathcal{V})$ else

	e_1				
e_2	\exists	T	F	U	-
	U	F	F	F	F
	-	-	F	U	-

$remove(\exists nseq \in e_1' : 'e_2, \mathcal{V}) =$
false iff $undefined(e_1, \mathcal{V}) \vee false(e_2, \mathcal{V})$
undefined iff $\neg undefined(e_1, \mathcal{V}) \wedge undefined(e_2, \mathcal{V})$
 $\forall nseq \in remove(e_1, \mathcal{V})' : 'remove(e_2, \mathcal{V})$ else

	e_1				
e_2	\exists_1	T	F	U	-
	U	F	F	F	F
	-	-	F	U	-

$remove(\exists_1 nseq \in e_1' : 'e_2, \mathcal{V}) =$
false iff $undefined(e_1, \mathcal{V}) \vee false(e_2, \mathcal{V})$
undefined iff $\neg undefined(e_1, \mathcal{V}) \wedge undefined(e_2, \mathcal{V})$
 $\forall nseq \in remove(e_1, \mathcal{V})' : 'remove(e_2, \mathcal{V})$ else

	e_2				
e_1	\leftrightarrow	T	F	U	-
	T	T	F	F	e_1
	F	F	T	F	$\neg e_1$
	U	F	F	U	F
	-	e_2	$\neg e_2$	F	$e_1 \leftrightarrow e_2$

Relational Operators

Binary relational operators form boolean expressions, comparing two subexpressions. Unlike boolean or arithmetical operators, it is not possible to omit the operator and retain one of the subexpressions, since the subexpressions are not boolean expressions. In our approach, we do not evaluate the relational operators $>$, $<$, \geq , \leq in respect to their truth-value. Therefore, these expressions are undefined if one of their subexpressions is undefined. Removal is defined accordingly.

		$op \in \{<, >, \leq, \geq\}$	
e_1	e_2	U	-
	U	U	U
	-	U	-

$remove(e_1 > e_2, \mathcal{V}) =$	undefined	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
	$remove(e_1, \mathcal{V}) > remove(e_2, \mathcal{V})$		else
$remove(e_1 < e_2, \mathcal{V}) =$	undefined	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
	$remove(e_1, \mathcal{V}) < remove(e_2, \mathcal{V})$		else
$remove(e_1 \geq e_2, \mathcal{V}) =$	undefined	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
	$remove(e_1, \mathcal{V}) \geq remove(e_2, \mathcal{V})$		else
$remove(e_1 \leq e_2, \mathcal{V}) =$	undefined	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
	$remove(e_1, \mathcal{V}) \leq remove(e_2, \mathcal{V})$		else

A special relational operator is the element-of operator $e_1 \in e_2$, where e_1 denotes an element and e_2 denotes a set. The element-of operator appears frequently in the guard of **if**-rules. The expression e_2 , denoting a set, is interpreted as the empty set if *undefined* holds. Therefore, *false* (*true*) holds for the element-of (not element-of) expression if e_2 is undefined. Likewise, an undefined expression should not be an element of any set. Note that according to this definition, *undefined* can not hold for an element-of expression.

		e_2	
e_1	\in	U	-
	U	F	F
	-	F	-

		e_2	
e_1	\notin	U	-
	U	T	T
	-	T	-

The equality operator is as significant as the element-of operator. For the equality operator, we take three special ASM elements into account - the element **undefined**, the empty set (\emptyset) and the empty sequence (*empty*). We interpretate an undefined expression e as **undefined**, empty set or empty sequence, depending on the context. Therefore, *true* holds if an undefined e is equated with one of these elements. Likewise,

false holds if an undefined expression is said to be unequal to one of these elements. Note that equality is symmetric, so if *true* holds for $e = \mathbf{undefined}$, it also holds for $\mathbf{undefined} = e$.

Excluding the cases addressed above, two expressions should never be equal if one expression is undefined and the other expression is not.

$$\begin{array}{ll}
\mathit{true}(e = \mathbf{undefined}, \mathcal{V}) & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
\mathit{false}(e \neq \mathbf{undefined}, \mathcal{V}) & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
\mathit{true}(e = \emptyset, \mathcal{V}) & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
\mathit{false}(e \neq \emptyset, \mathcal{V}) & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
\mathit{true}(e = \mathit{empty}, \mathcal{V}) & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
\mathit{false}(e \neq \mathit{empty}, \mathcal{V}) & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
\mathit{false}(e_1 = e_2, \mathcal{V}) & \text{iff } \neg \mathit{true}(e_1 = e_2, \mathcal{V}) \wedge \\
& (\mathit{undefined}(e_1, \mathcal{V}) \wedge \neg \mathit{undefined}(e_2, \mathcal{V}) \vee \\
& \neg \mathit{undefined}(e_1, \mathcal{V}) \wedge \mathit{undefined}(e_2, \mathcal{V})) \\
\mathit{true}(e_1 \neq e_2, \mathcal{V}) & \text{iff } \neg \mathit{false}(e_1 \neq e_2, \mathcal{V}) \wedge \\
& (\mathit{undefined}(e_1, \mathcal{V}) \wedge \neg \mathit{undefined}(e_2, \mathcal{V}) \vee \\
& \neg \mathit{undefined}(e_1, \mathcal{V}) \wedge \mathit{undefined}(e_2, \mathcal{V})) \\
\mathit{undefined}(e_1 = e_2, \mathcal{V}) & \text{iff } \mathit{undefined}(e_1, \mathcal{V}) \wedge \mathit{undefined}(e_2, \mathcal{V}) \\
\mathit{undefined}(e_1 \neq e_2, \mathcal{V}) & \text{iff } \mathit{undefined}(e_1, \mathcal{V}) \wedge \mathit{undefined}(e_2, \mathcal{V})
\end{array}$$

Arithmetic Operators

Arithmetic operators are removed if one of their subexpressions is undefined. Removal is less strict for plus and minus, which are only removed if both expressions are undefined.

$$\begin{array}{ll}
\mathit{remove}(-e, \mathcal{V}) = & \\
\mathbf{undefined} & \text{iff } \mathit{undefined}(e, \mathcal{V}) \\
-\mathit{remove}(e, \mathcal{V}) & \text{else} \\
\mathit{remove}(e_1 + e_2, \mathcal{V}) = & \\
\mathbf{undefined} & \text{iff } \mathit{undefined}(e_1, \mathcal{V}) \wedge \mathit{undefined}(e_2, \mathcal{V}) \\
\mathit{remove}(e_1, \mathcal{V}) & \text{iff } \mathit{undefined}(e_2, \mathcal{V}) \\
\mathit{remove}(e_2, \mathcal{V}) & \text{iff } \mathit{undefined}(e_1, \mathcal{V}) \\
\mathit{remove}(e_1, \mathcal{V}) + \mathit{remove}(e_2, \mathcal{V}) & \text{else} \\
\mathit{remove}(e_1 - e_2, \mathcal{V}) = & \\
\mathbf{undefined} & \text{iff } \mathit{undefined}(e_1, \mathcal{V}) \wedge \mathit{undefined}(e_2, \mathcal{V}) \\
\mathit{remove}(e_1, \mathcal{V}) & \text{iff } \mathit{undefined}(e_2, \mathcal{V}) \\
\mathit{remove}(e_2, \mathcal{V}) & \text{iff } \mathit{undefined}(e_1, \mathcal{V}) \\
\mathit{remove}(e_1, \mathcal{V}) - \mathit{remove}(e_2, \mathcal{V}) & \text{else} \\
\mathit{remove}(e_1 * e_2, \mathcal{V}) = &
\end{array}$$

undefined $remove(e_1, \mathcal{V}) * remove(e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$ else
$remove(e_1 / e_2, \mathcal{V}) =$ undefined $remove(e_1, \mathcal{V}) / remove(e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$ else
$remove(e_1 \text{ MOD } e_2, \mathcal{V}) =$ undefined $remove(e_1, \mathcal{V}) \text{ MOD } remove(e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$ else
$remove(e_1 \text{ DIV } e_2, \mathcal{V}) =$ undefined $remove(e_1, \mathcal{V}) \text{ DIV } remove(e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$ else

$undefined(-e, \mathcal{V})$	iff	$undefined(e, \mathcal{V})$
$undefined(e_1 + e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \wedge undefined(e_2, \mathcal{V})$
$undefined(e_1 - e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \wedge undefined(e_2, \mathcal{V})$
$undefined(e_1 * e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
$undefined(e_1 / e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
$undefined(e_1 \text{ mod } e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$
$undefined(e_1 \text{ rem } e_2, \mathcal{V})$	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_2, \mathcal{V})$

Sets and Sequences

Sequence (set) composition constructs a sequence (set) by applying expression e_1 to elements from e_2 for which the boolean expression e_3 holds. If expression e_1 or e_3 are undefined, so is the sequence/set. In that case, *false* will hold when testing if an element is contained in the structure. If *undefined* holds for e_2 (i.e., e_2 is interpreted as empty) or *false* holds for e_3 , the resulting sequence is empty, and the resulting set is the empty set.

$$\langle e_1 \mid x \in e_2 : e_3 \rangle, \{e_1 \mid x \in e_2 : e_3\}$$

$remove(\langle e_1 \mid x \in e_2 : e_3 \rangle, \mathcal{V}) =$ undefined <i>empty</i>	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_3, \mathcal{V})$
$\langle remove(e_1, \mathcal{V}) \mid x \in remove(e_2, \mathcal{V}) : remove(e_3, \mathcal{V}) \rangle$	iff	$(false(e_3, \mathcal{V}) \vee undefined(e_2, \mathcal{V}))$ $\wedge \neg undefined(e_1, \mathcal{V})$
else		
$remove(\{e_1 \mid x \in e_2 : e_3\}, \mathcal{V}) =$ undefined \emptyset	iff	$undefined(e_1, \mathcal{V}) \vee undefined(e_3, \mathcal{V})$
	iff	$(false(e_3, \mathcal{V}) \vee undefined(e_2, \mathcal{V}))$ $\wedge \neg undefined(e_1, \mathcal{V})$

$$\{ \text{remove}(e_1, \mathcal{V}) \mid x \in \text{remove}(e_2, \mathcal{V}) : \text{remove}(e_3, \mathcal{V}) \} \quad \text{else}$$

$$\begin{aligned} \text{undefined}(\langle e_1 \mid x \in e_2 : e_3 \rangle, \mathcal{V}) & \text{ iff } \text{undefined}(e_1, \mathcal{V}) \vee \text{undefined}(e_3, \mathcal{V}) \\ \text{undefined}(\{e_1 \mid x \in e_2 : e_3\}, \mathcal{V}) & \text{ iff } \text{undefined}(e_1, \mathcal{V}) \vee \text{undefined}(e_3, \mathcal{V}) \end{aligned}$$

Operations on sets and sequences need at least one argument for which *undefined* doesn't hold, otherwise *undefined* holds for the operation. Removal is straightforward, the definition is equivalent to the definition for the arithmetic operators plus and minus.

e_1	op	op $\in \{ , \cup \}$	\cup	-
		\cup	-	-
e_1	op	op $\in \{ \cap, \cup, \cap \}$	\cup	-
		\cup	-	-
		-	-	-

For the operations map (\mapsto) and range (\dots), *undefined* holds if *undefined* holds for any subexpression. The definition of removal is equivalent to the arithmetic operator *.

e_1	op	op $\in \{ \mapsto, \dots \}$	\cup	-
		\cup	\cup	\cup
		-	\cup	-

Function and Macro Calls

A macro call is removed if one of the parameters passed is undefined, or the macro name itself has been marked as undefined. Otherwise, parameters that correspond to removed formal parameters from the rule macro definition are removed from the parameter list of the macro call. This is done with the function *remfpar* defined in Section 4.3.2. *count* is a natural number that holds the information which formal parameters were removed.

$$\begin{aligned} \text{remove}(\underline{\text{MkName}}(), \mathcal{V}) & = \text{undefined}(\underline{\text{MkName}}()) \quad \text{iff } \text{undefined}(\text{MkName}) \\ & \quad \text{else} \\ \text{remove}(\underline{\text{MkName}}(\underline{fcs}), \mathcal{V}) & = \text{skip} \quad \text{iff } \text{undefined}(\text{MkName}) \vee \text{undefined}(fcs, \mathcal{V}) \\ & \quad \text{MkName}(\underline{\text{removepar}}(fcs, \text{numfpar}(fcs) - 1, \text{count}(\text{MkName}))) \\ & \quad \text{else} \end{aligned}$$

$$\begin{aligned} \mathit{undefined}(\mathit{MkName}(_), \mathcal{V}) & \text{ iff } \mathit{undefined}(\mathit{MkName}) \\ \mathit{undefined}(\mathit{MkName}(fcs)_), \mathcal{V}) & \text{ iff } \mathit{undefined}(\mathit{MkName}) \vee \mathit{undefined}(fcs, \mathcal{V}) \end{aligned}$$

Function calls are similar to rule macro calls. A function call can either refer to a location of the ASM, or a derived function defining an expression. Removal for function calls is identical to removal for macro calls, removing the call if either the parameters or the function itself are undefined.

$$\begin{aligned} \mathit{remove}(f, \mathcal{V}) & = \\ & \quad \mathit{undefined} & \text{ iff } \mathit{undefined}(f) \\ & \quad f & \text{ else} \\ \mathit{remove}(exp.f, \mathcal{V}) & = \\ & \quad \mathit{undefined} & \text{ iff } \mathit{undefined}(f) \vee \mathit{undefined}(exp, \mathcal{V}) \\ & \quad \mathit{remove}(exp, \mathcal{V}).f & \text{ else} \\ \mathit{remove}(f(fcs)_), \mathcal{V}) & = \\ & \quad \mathit{undefined} & \text{ iff } \mathit{undefined}(f) \vee \mathit{undefined}(fcs, \mathcal{V}) \\ & \quad f(_ \mathit{removepar}(fcs, \mathit{numfpar}(fcs) - 1, \mathit{count}(f))) & \text{ else} \end{aligned}$$

$$\begin{aligned} \mathit{undefined}(f(fcs)_), \mathcal{V}) & \text{ iff } \mathit{undefined}(f) \vee \mathit{undefined}(fcs, \mathcal{V}) \\ \mathit{undefined}(f, \mathcal{V}) & \text{ iff } \mathit{undefined}(f) \\ \mathit{undefined}(exp.f, \mathcal{V}) & \text{ iff } \mathit{undefined}(f) \vee \mathit{undefined}(exp, \mathcal{V}) \end{aligned}$$

Names and keywords in expressions are removed if they are marked as undefined, otherwise they are left untouched. Predicate *undefined* never holds for literals (for example, natural numbers) and the special ASM element **undefined**.

$$\begin{aligned} \mathit{remove}(\mathit{DomainName}, \mathcal{V}) & = \mathbf{if} \mathit{undefined}(\mathit{DomainName}) \mathbf{then} \mathit{undefined} \mathbf{else} \mathit{DomainName} \mathbf{endif} \\ \mathit{remove}(\mathit{SynName}, \mathcal{V}) & = \mathbf{if} \mathit{undefined}(\mathit{SynName}) \mathbf{then} \mathit{undefined} \mathbf{else} \mathit{SynName} \mathbf{endif} \\ \mathit{remove}(\mathit{ProgramName}, \mathcal{V}) & = \mathbf{if} \mathit{undefined}(\mathit{ProgramName}) \mathbf{then} \mathit{undefined} \mathbf{else} \mathit{ProgramName} \mathbf{endif} \\ \mathit{remove}(\mathit{Keyword}, \mathcal{V}) & = \mathbf{if} \mathit{undefined}(\mathit{Keyword}) \mathbf{then} \mathit{undefined} \mathbf{else} \mathit{Keyword} \mathbf{endif} \\ \mathit{remove}(\mathit{Literal}, \mathcal{V}) & = \mathit{Literal} \\ \mathit{remove}(\mathit{undefined}, \mathcal{V}) & = \mathit{undefined} \\ \mathit{remove}(\leq pcs \geq, \mathcal{V}) & = \leq \mathit{remove}(pcs, \mathcal{V}) \geq \end{aligned}$$

$$\begin{aligned} \mathit{undefined}(\mathit{DomainName}, \mathcal{V}) & \text{ iff } \mathit{undefined}(\mathit{DomainName}) \\ \mathit{undefined}(\mathit{SynName}, \mathcal{V}) & \text{ iff } \mathit{undefined}(\mathit{SynName}) \\ \mathit{undefined}(\mathit{ProgramName}, \mathcal{V}) & \text{ iff } \mathit{undefined}(\mathit{ProgramName}) \\ \mathit{undefined}(\mathit{Keyword}, \mathcal{V}) & \text{ iff } \mathit{undefined}(\mathit{Keyword}) \end{aligned}$$

$undefined(\text{Literal}, \mathcal{V})$ iff **false**
 $undefined(\text{undefined}, \mathcal{V})$ iff **false**
 $undefined(\leq \text{pcs} \geq, \mathcal{V})$ iff $undefined(\text{pcs}, \mathcal{V})$

5 Consistency of SDL Profiles

A set of SDL profiles is called *consistent*, if any specification that can be stated in these profiles behaves exactly the same way in each profile. Deriving the profiles from a common language definition enables us to make statements about consistency, because, unlike profiles defined independently, the derived profiles share many common parts.

A run of an ASM is a sequence of states, where each subsequent state is the result of firing all rules whose conditions are true on the preceding state. For non-deterministic, multi-agent ASMs, the legal behaviour is given by a set of runs, each run in the set describing a possible execution of the system. Two SDL profiles are considered consistent, if they yield the same set of runs of their respective ASMs for all specifications contained in both profiles.

In order to prove consistency, it is sufficient to show that only dead ASM rules are removed. This property does not follow automatically from the formally defined operations for removal, since they rely on heuristics in some parts. However, based on these operations, it is possible to derive proof obligations that have to be verified in order to prove consistency.

E.g., during removal, an **if**-rule can be replaced by the subrule in the **then**-block of the rule, if the predicate *true* holds for the guard. To prove consistency, it is sufficient to prove that for all specifications of the SDL profile, the guard evaluates to **true** in all reachable states¹. Likewise, if the predicate *false* holds for the guard, we have to prove that for all specifications of the SDL profile, the guard evaluates to **false** in all reachable states. In case *undefined* holds for the guard, we have to prove that the **if**-statement cannot be reached at all.

ASM-Listing 1 shows a part of the formal language definition that was removed as part of the save feature of SDL. For SDL profiles that do not contain the save feature, no grammatical elements of *Save-signalset* exist. Therefore, selecting the *Save-signalset* for any state yields **undefined**, and selecting *Signal-identifier-set* for the element **undefined** yields the empty set. Since *Save-signalset* is not modified in the formal language definition, this holds for any reachable state of the ASM. An element cannot be contained in an empty set, therefore the guard is always false, and omitting the **if**-statement leads to a consistent definition for specifications without save.

```
1 if Self.signalChecked.signalType ∈  
2   sn.stateAS1.s-Save-signalset.s-Signal-identifier-set then  
3   Self.SignalSaved := True
```

¹This condition is very generic. It would suffice to show that the guard is always **true** for all reachable states that lead to the firing of the **if**-rule.

Choose. Choose nondeterministically selects an element that satisfies the constraint given by expression exp . If a **choose**-rule is removed, we have to prove consistency by proving the expression exp to be **false** in any reachable state, and therefore - according to the semantics of ASMs - the **choose**-rule equates to an empty update set. Alternatively, we can prove that the **choose**-rule cannot be reached.

Extend. **extend**-rules are removed if they extend a domain that has been removed from the ASM signature. For a domain that is associated with a language feature, an extension of the domain must not be reached if that feature is removed. In order to prove consistency, we therefore have to prove that such a rule cannot be reached.

Rule Macro Definition. A rule macro can be removed without affecting consistency if no corresponding rule call exists in a reachable part of the ASM, or if the body of the rule macro can be reduced to skip while maintaining consistency.

Boolean Expressions Parts of boolean expressions are removed if they have no influence on the final result, e.g. if *true* holds for a subexpression of a conjunction. In this case, the proof obligation is to show that the subexpression is always true for specifications of the SDL profile.

Proof obligations on boolean expressions can be split into proof obligations on subexpressions, as shown for \wedge and \vee below. E.g., in order to prove consistency for predicate *true* on $e_1 \wedge e_2$, we can prove consistency for predicate *true* on e_1 and e_2 .

$$true(e_1 \wedge e_2) \text{ iff } true(e_1) \textbf{ and } true(e_2) \quad (5.1)$$

$$false(e_1 \wedge e_2) \text{ iff } false(e_2) \textbf{ or } false(e_1) \quad (5.2)$$

$$true(e_1 \vee e_2) \text{ iff } true(e_1) \textbf{ or } true(e_2) \quad (5.3)$$

$$false(e_1 \vee e_2) \text{ iff } false(e_1) \textbf{ and } false(e_2) \quad (5.4)$$

Relational Operators Relational Operators like $<$, $>$, \leq , \geq are undefined if one of their subexpressions is undefined. To prove consistency, in this case we have to prove that the expression cannot be reached.

In case of the element-of operator $e_1 \in e_2$, two heuristics were used. If *undefined* holds for the expression on the right hand side, the set is interpreted as empty and *true* holds for the expression. In this case, we have to prove that the right hand side always equates to the empty set. In case *undefined* holds for the expression on the left hand side, we have to prove that the element described by this expression is not contained in the set on the right hand side in any reachable state.

6 SDL-Profile Tool

Based on the formalisation provided in Section 4, we have implemented a tool called SDL-profile tool in order to automate the reduction process, providing visible results. The tool reads the formal semantics definition, performs the *remove* operation based on a *reduction profile*, and outputs a reduced version of the formal semantics. The reduction profile is a list of domain names, function names and macro names that are removed from the ASM signature (or from the set of rules, in the case of macro names), possibly defining default values, e.g. for predicates. Figure 6.1 shows the sequence of steps performed during the removal, and the tools used for each step.

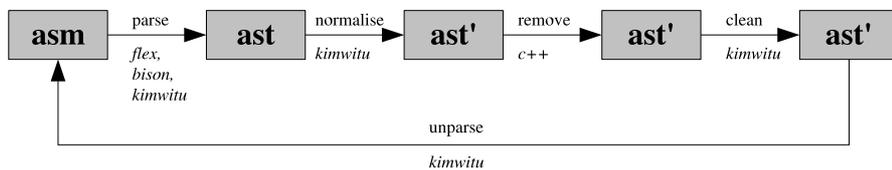


Figure 6.1: Tool Chain of the SDL-Profile Tool

6.1 Tool Chain

Parser

The *parser* takes an ASM specification as input and creates an abstract syntax tree representation of the specification as output. It is generated out of specifications of the lexis, grammar and abstract syntax of Abstract State Machines, as used in the formal semantics of SDL-2000 [3]. The specification of the abstract syntax is translated by *kimwitu++* [12] to a data structure for the abstract syntax tree, using C++ classes. Scanner and parser are generated by *flex* and *bison*, respectively. Apart from minor differences, the parser is identical to the parser used in [16].

Normalisation

The *normalisation* step transforms the abstract syntax tree to a pre-removal normal form. The transformation is specified by rewrite rules on the abstract syntax tree. The rewrite rules are translated to C++ functions by the *kimwitu* tool. The main function of the normalisation step is to split up complicated abstract syntax rules, in order to make the definition of the remove function easier.

Remove

The *remove* step is the implementation of the removal formalised in Section 4. For each type of node (*phylum*) in the abstract syntax definition, a remove function is introduced. The remove function performs removal for each term of the respective phylum, e.g. the terms **IfThenElse**, **Choose**, and **Extend** for the *rule* phylum. It returns a term of the respective phylum as result – e.g. the remove function for rules always returns a term of type rule.

For a term of a phylum, removal starts by checking conditions consisting of the predicates *true*, *false* and *undefined*, as defined in the formalisation of the reduction process. If a condition evaluates to **true**, a modified term is returned, calling remove recursively on the subterms of the term if necessary. E.g., for the rule term **IfThenElse**, if the predicate *true* holds for expression *exp*, removal continues with the **then**-part, if the predicate *false* holds for expression *exp*, removal continues with the **else**-part. If *undefined* holds for the expression *exp*, the rule term **Skip** is returned.

```
IfThenElse(exp, r1, r2): {
  if (eval_true(exp,V)) { return remove(r1,V); };
  if (eval_false(exp,V)) { return remove(r2,V); };
  if (eval_undef(exp,V)) { return Skip(); };
  return IfThenElse(remove(exp,V), remove(r1,V), remove(r2,V));
}
```

Cleanup

The *cleanup* step transforms superfluous rules resulting from the removal step to a post-removal normal form. The normal form is achieved by defining term rewrite rules in kimwitu. Unlike removal, the rewrite rules apply anywhere where their left hand side matches, and are applied as long as a match is found.

The cleanup step only removes trivial parts of the ASM specification. The resulting specification is semantically equivalent to the specification before the cleanup step.

Iteration

Given a completely defined reduction profile, only one run of the SDL-profile tool is needed to generate a reduced ASM semantics definition. In case the reduction profile is incomplete, the profile tool can identify further names in the ASM signature that can be removed, and iterate the removal process. E.g., a function in the ASM signature with a target domain that has been removed during the previous removal step, is included in the reduction profile of a subsequent iteration.

Unparsing

Unparsing traverses the abstract syntax tree and outputs a string representation of every node. The result is a textual representation of the formal semantics tree in the original input format. Therefore, the output of the profile tool can be used as the input

for a subsequent run of the profile tool. We also provide different output formats, e.g. a Latex document of the formal semantics, or a compilation to C++.

6.2 Application of the SDL-Profile Tool

Given an ASM formal semantics definition and a reduction profile, the SDL-profile tool generates a reduced formal semantics definition in the original format. In order to validate the removal process, we compare the original semantics definition with the reduced version. For this, we have used graphical diff-based tools (e.g., tkdiff) to highlight the differences between the versions. Using the SDL-profile tool, we have created reduction profiles for several language modules, such as timers, exceptions, save, composite states and inheritance. We have also created reduction profiles for language profiles like *SAFIRE* and *Core*, resulting in formal semantics definitions that, with small modifications, match these SDL profiles.

ASM Listings 1 and 2 show an excerpt of the formal semantics definition before and after applying the SDL-profile tool, using a reduction profile for SDL exceptions. The reduction profile contains, besides other function and macro names, the function name *currentExceptionInst*, which is interpreted as *undefined* in the context below. Therefore, the predicate *false* holds for the guard of the **if**-rule, and the first part of the **if**-rule is removed.

```

1 SELECTTRANSITIONSTARTPHASE ≡
2   if (Self.currentExceptionInst ≠ undefined) then
3     Self.agentMode3 := selectException
4     Self.agentMode4 := startPhase
5   elseif (Self.currentStartNodes ≠  $\emptyset$ ) then
6     ...
7   else
8     ...
9   endif

```

ASM-Listing 1: Macro SELECTTRANSITIONSTARTPHASE before Reduction

```

1 SELECTTRANSITIONSTARTPHASE ≡
2   if (Self.currentStartNodes ≠  $\emptyset$ ) then
3     ...
4   else
5     ...
6   endif

```

ASM-Listing 2: Macro SELECTTRANSITIONSTARTPHASE after Reduction

7 Related Work

A modular language definition as described in this paper can be found in the language specification of UML [15, 14]. The abstract syntax of UML is defined using a meta-model approach, using classes to define language elements and packages to group language elements into medium-grained units. The core of the language is defined by the Kernel package, specifying basic elements of the language such as packages, classes, associations and types. However, each meta-model class/language element has only an informal description of its semantics.

UML has a profile mechanism that allows metaclasses from existing metamodels to be extended and adapted, using stereotypes. Semantics and constraints may be added as long as they are not in conflict with the existing semantics and constraints. E.g., the profile mechanism has been used to define a UML profile for SDL, enabling the use of UML 2.0 as a front-end for SDL-2000 [11].

ConTraST [1] is an SDL to C++ transpiler that generates a readable C++ representation of an SDL specification by preserving as much of the original structure as possible. The generated C++ code is compiled together with a runtime environment that is a C++ implementation of the formal semantics defined in Z100.F3. ConTraST is based on the textual syntax of SDL-96, and supports SDL profiles syntactically through deactivation of language features.

In [13], the concept of program slicing is extended to Abstract State Machines. For an expressive class of ASMs, an algorithm for the computation of a minimal slice of an ASM, given a slicing criterion, is presented. While the complexity of the algorithm is acceptable in the average case, the worst case complexity is exponential.

8 Conclusions and Outlook

In this paper, we have introduced the concept of SDL profiles as well-defined subsets of SDL, leading to smaller, more understandable language definitions. Tool support can be based on these profiles, leading to faster tool development and less expensive tools. Based on the smaller language definitions, code optimisations can be performed when generating code from a specification.

We have argued for the importance of formal semantics for language definitions, and the importance of deriving the formal semantics of SDL profiles from a common formal semantics definition. This allows us to compare the formal semantics of different SDL profiles, and to make assertions about their consistency.

To achieve deterministic results, we have formalised the extraction of the formal semantics for SDL profiles from the complete formal semantics of SDL-2000. The extraction is based on recognising and removing dead ASM rules from the formal semantics definition, starting from a reduced ASM signature. The reduction of the ASM signature is derived from the abstract syntax of removed language modules. The extraction has been automated by the SDL-profile tool, providing visible results. This tool has been used to create several language profiles for SDL-2000, by removing SDL language modules from the formal semantics definition, such as e.g. exceptions, timers, save and composite states. The reduction achieved is significant. E.g., the formal semantics definition for SAFIRE has been reduced to less than 2300 lines of specification, and less than 1100 lines for a small core of SDL, from about 3700 lines of the complete formal semantics.

Based on the formally defined process for the derivation of formal language definitions for SDL profiles, we can define precise criteria for the consistency of SDL profiles. Currently, some consistency criteria have to be verified manually. Our future work will focus on improving the extraction process, so that further criteria can be checked automatically.

Bibliography

- [1] FLIEGE, Ingmar ; GRAMMES, Rüdiger ; WEBER, Christian: ConTraST - A Configurable SDL Transpiler And Runtime Environment. In: GOTZHEIN, Reinhard (Hrsg.): *SAM'06 - Fifth Workshop on System Analysis and Modelling, Kaiserslautern, Germany*, 2006. – to be published
- [2] GLÄSSER, Uwe ; GOTZHEIN, Reinhard ; PRINZ, Andreas: The Formal Semantics of SDL-2000 - Status and Perspectives. In: *Computer Networks* 42 (2003), Nr. 3, S. 343–358
- [3] GLÄSSER, Uwe ; GOTZHEIN, Reinhard ; PRINZ, Andreas: An Introduction To Abstract State Machines / Department of Computer Science, University of Kaiserslautern. 2003 (326/03). – Forschungsbericht
- [4] GUREVICH, Yuri: Evolving Algebras 1993: Lipari Guide. In: BÖRGER, Egon (Hrsg.): *Specification and Validation Methods*. Oxford University Press, 1995, S. 9–36
- [5] GUREVICH, Yuri: May 1997 Draft of the ASM Guide / EECS Department, University of Michigan. 1997 (CSE-TR-336-97). – Forschungsbericht
- [6] ITU: *Recommendation Z.100 (08/02): Specification and Description Language (SDL)*. Geneva, 2002
- [7] ITU: *Recommendation Z.100 (2002) Amendment 1 (10/03): Specification and Description Language (SDL)*. Geneva, 2003
- [8] ITU: *Recommendation Z.100 (2002) Corrigendum 1 (08/04): Specification and Description Language (SDL)*. Geneva, 2004
- [9] ITU STUDY GROUP 10: *Draft Z.100 Annex F2 (11/00)*. 2000
- [10] ITU STUDY GROUP 10: *Draft Z.100 Annex F3 (11/00)*. 2000
- [11] ITU STUDY GROUP 17: *UML Profile for SDL*. 2005. – Draft Recommendation Z.109
- [12] LÖWIS, Martin von ; PIEFEL, Michael: The Term Processor Kimwitu++. In: CALLAOS, Nagib (Hrsg.) ; HERNÁNDEZ-ENCINAS, Luis (Hrsg.) ; YETIM, Fahri (Hrsg.): *SCI 2002: The 6th World Multiconference on Systemics, Cybernetics and Informatics, Orlando, USA*, 2002, S. 182–186

- [13] NOWACK, Antje: Slicing Abstract State Machines. In: ZIMMERMANN, Wolf (Hrsg.) ; THALHEIM, Bernhard (Hrsg.): *Abstract State Machines 2004. Advances in Theory and Practice, Lutherstadt Wittenberg, Germany* Bd. 3052, Springer, Januar 2004 (LNCS), S. 186–201
- [14] OBJECT MANAGEMENT GROUP: *Unified Modeling Language Specification, Version 1.3*. 2000. – www.uml.org
- [15] OBJECT MANAGEMENT GROUP: *Unified Modeling Language: Superstructure, Version 2.0*. 2005. – www.uml.org
- [16] PRINZ, Andreas ; LÖWIS, Martin von: Generating a Compiler for SDL from the Formal Language Definition. In: REED, Rick (Hrsg.) ; REED, Jeanne (Hrsg.): *SDL 2003: System Design* Bd. 2708, Springer, 2003 (LNCS), S. 150–165
- [17] SDL TASK FORCE: *SDL+ - The Simplest, Useful 'Enhanced SDL-Subset' for the Implementation and Testing of State Machines*. 2005. – www.sdltaskforce.org