SMT for state-based formal methods: the ASM case study

Paolo Arcaini*
Charles University
Faculty of Mathematics and Physics,
Czech Republic
arcaini@d3s.mff.cuni.cz

Angelo Gargantini
DIGIP
University of Bergamo, Italy
angelo.gargantini@unibg.it

Elvinia Riccobene
Department of Computer Science
University of Milan, Italy
elvinia.riccobene@unimi.it

ABSTRACT

State-based transition systems can take advantage of a symbolic representation of the concepts of state and transition in order to automatically solve verification questions that could not be otherwise tackled in terms of explicit representation of the transition system. We report here our experience in developing solutions, approaches and supporting tools of verification problems regarding the Abstract State Machines (ASMs), a transition system which can be considered as an extension of Finite State Machines. We present the symbolic representation of an ASM and of its computational model in terms of the Yices SMT solver. We also discuss two scenarios of verification questions regarding the ASMs for which the symbolic representation helped us to formalize and solve the problem by satisfiability checking, namely automatic proof of correct ASM refinement and runtime verification.

CCS CONCEPTS

• **Software and its engineering → Formal methods**; **Automated static analysis**;

KEYWORDS

Abstract State Machines, SMT, Yices, refinement proof, runtime verification

ACM Reference format:
Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2017. SMT for state-based formal methods: the ASM case study. In *Proceedings of ACM conference, NASA Ames Research Center, USA, May 2017 (AFM 2017),* 9 pages.
https://doi.org/10.475/123_4

1 INTRODUCTION

The architectural and behavioural complexity of modern systems and the need to guarantee critical properties yet at the early stages of the system life-cycle, require a rigorous development process based on the use of formal methods for system specification, validation and verification. Formal models can be used to understand if the system under development satisfies the given requirements (validation) and guarantees system properties (verification).

Specification can be given in operational style, when the system behaviour is expressed in terms of states and computation steps, or in declarative style, when the system is specified in terms of holding properties. There has been an endless debate about which style fits better the designer needs: some argue that with an operational style designers tend to insert implementation details in the abstract specifications, others observe that practitioners feel uncomfortable with declarative notations like temporal logics. Operational specifications are easier to write and understand; however, declarative specifications could be more suitable for verification purposes.

Abstract State Machines (ASMs) [14] are a state-based operational formal method that has been widely used as a system engineering method in different contexts: definition of industrial standards for programming and modeling languages, design of industrial control systems, modeling service and cloud systems, design and analysis of protocols, architectural design, language design, verification of compilers, etc.

Originally defined by Y. Gurevich in 1993 with the goal to sharpen the Church-Turing thesis [18], along the years, ASMs have been used in the field of software engineering to model systems and investigate their properties. To this aim, the usage of such formalism provides benefits under several viewpoints. If we consider *expressive power*, ASMs represent a general model of computation where the static view of a system is represented by means of a mathematical algebra and its dynamics is given in terms of transition rules. Concerning *understandability*, ASMs provide a way to describe behavioral issues by means of pseudo-code working over abstract data structures; therefore, ASM models are easily understandable without strong mathematical skills. If we consider *methodological issues*, the ASM formalism is the basis of a rigorous development method based on the concept of a "ground model" representing a precise but concise high-level formalization of the system, and on the "refinement principle" that allows to capture all details of the system design by a sequence of refined models till the desired level of detail. For *analysis purposes*, the rigorous mathematical foundation of ASM models, allows the application of formal techniques for model validation and verification. From the *implementation point of view*, the simple pseudo-code can be translated into a high level programming language source code in a quite simple manner (even automatically [12]).

To facilitate the practical usage of ASMs as software engineering formal method and to overcome the lack of automated tool support and poor tools integration around ASMs, in the past we worked on the development of the ASMETA framework [7], allowing ASMs to be used in an efficient and tool supported manner during the entire software development life cycle. In [6, 8], we have also indicated the process usually followed to develop systems from the definition of informal requirements to code implementation and execution.

*The research reported in this paper has been partially supported by the Czech Science Foundation project number 17-12465S.

AFM 2017, May 2017, NASA Ames Research Center, USA

ACM ISBN 123-4567-24-567/08/06…$15.00
https://doi.org/10.475/123_4

It was by developing specific model analysis approaches around ASMs that we came across the necessity of having a symbolic representation of an ASM and of its computational model. Although a representation of ASMs in the theorem prover PVS already exists [17], it requires some expertise in theorem proving in order to be used. Therefore, we preferred to explore the use of Satisfiability Modulo Theories (SMT) solvers that guarantee a higher degree of automation. Thanks this representation in SMT, we were able to reduce verification questions, that would have been unfeasible to solve automatically by means of explicit representation of ASM states and of ASM computations as explicit sequences of states, to SMT problems. In particular, proving correctness of ASM model refinement [5] and runtime verification through nondeterministic ASMs [4] were solved as SMT satisfiability checking. In both cases, the symbolic representation helps *to formalize the problem* in a very concise way by means of a *logical context* representing ASM states and computation steps, and *to solve the problem* as satisfability checking of the context.

In this position paper, we present the symbolic representation of an ASM and of its computational model in terms of a context of the Yices SMT solver. We then introduce the formalization of the two main ASM verification questions we tackled by the use of Yices: proof of the correct ASM refinement and runtime verification. Along the presentation, we discuss the problems we faced to move from an explicit representation of the machine state and its computation step to a symbolic one. We also motivate formalization choices made for reducing a computational problem to a satisfiability problem.

The paper is organized as follows. Sect. 2 briefly introduces the Abstract State Machines. Sect. 3 presents the symbolic representation in Yices of an ASM and how to build an SMT context representing the initial state and the computation step from a given ASM model. Sect. 4 discusses the problem of reducing to an SMT problem satisfaction the question of verifying correctness of an ASM model refinement step. Sect. 5 introduces the problem of runtime verification using ASM nondeterministic models and shows how to formalize its solution in terms of a symbolic SMT problem. We conclude the paper with Sect. 6, where we discuss strengths and weaknesses of our symbolic representation of the ASMs and we outline possible further verification contexts where this representation could be exploited.

2 ABSTRACT STATE MACHINES

Abstract State Machines (ASMs) [14] are an extension of FSMs, where unstructured control states are replaced by states with arbitrarily complex data.

ASM *states* are algebraic structures, i.e., domains of objects with functions and predicates defined on them. An ASM *location*, defined as the pair (*function-name*, *list-of-parameter-values*), represents the abstract ASM concept of basic object containers. The couple (*location*, *value*) represents a machine memory unit. Therefore, ASM states can be viewed as abstract memories.

Location values are changed by firing *transition rules*. They express the modification of functions interpretation from one state to the next one. Note that the algebra signature is fixed and that functions are total (by interpreting undefined locations $f(x)$ with value *undef*). Location *updates* are given as assignments of the form $loc := v$, where *loc* is a location and v its new value. They are the basic units of rules construction. There is a limited but powerful set of *rule constructors* to express: guarded actions (`if-then`), simultaneous parallel actions (`par`), sequential actions (`seq`), nondeterminism (existential quantification `choose`), and unrestricted synchronous parallelism (universal quantification `forall`).

An ASM *computation* is, therefore, defined as a finite or infinite sequence $S_0, S_1, \ldots, S_n, \ldots$ of states of the machine, where S_0 is an initial state and each S_{n+1} is obtained from S_n by firing the unique *main rule* which in turn could fire other transitions rules. An ASM can have more than one *initial state*.

During a machine computation, not all the locations can be updated. Functions are classified as *static* (never change during any run of the machine) or *dynamic* (may change as a consequence of agent actions or *updates*). Dynamic functions are distinguished between *monitored* (only read by the machine and modified by the environment) and *controlled* (read and written by the machine). A further classification is between *basic* and *derived* functions, i.e., those coming with a specification or computation mechanism given in terms of other functions. It is possible to specify state *invariants*.

An ASM can be *nondeterministic* due to the presence of monitored functions (*external* nondeterminism) and of choose rules (*internal* nondeterminism).

A set of tools exists to support the ASM modeling process. Tools are part of the ASMETA (ASM mETAmodeling) framework[1] [7], and are strongly integrated in order to permit reusing information about models during different development phases. ASMETA provides basic functionalities for ASM model editing, and supports advanced model analysis techniques (as validation, verification, testing, model review, runtime verification, refinement proof, etc.).

3 ASM SYMBOLIC REPRESENTATION

An SMT problem is a decision problem for logical formulas with respect to combinations of background theories expressed in classical first-order logic with equality. An SMT instance is a generalization of a boolean SAT instance in which various sets of variables are replaced by predicates from a variety of underlying theories. SMT solvers can be used, as in our case, as automatic theorem provers by checking unsatisfiability. We use Yices [15] as SMT solver.

The following sections describe the mapping (in terms of *mapping functions*) from ASM models to Yices elements. The theories needed for the mapping are: uninterpreted functions, and linear and non-linear integer and real arithmetic.

The way to use the Yices elements obtained by such mapping depends on the technique that uses them; we will give examples of these usages in Sects. 4 and 5 in which we describe how we exploited the ASM symbolic representation for proving ASM refinement correctness and for runtime verification.

We first describe the mapping of the signature, of terms, of function definitions, and of transition rules. Then, we describe how to exactly capture the semantics of an ASM computation step, and we provide the complete description of the mapping of the initial state and of a generic step. We finally provide an example

[1] http://asmeta.sourceforge.net/

ASM domains declarations	**Yices**
enum domain D = {E_1, . . ., E_n}	(**define−type** D (scalar E_1. . .E_nD_undef))
Boolean	bool
Integer	int
Natural	nat

Table 1: T_{dom}: Mapping schema of ASM domains to Yices

ASM function declaration funcType ∈ {controlled, monitored, derived, static}	**Yices**
funcType f: Dom	(**define** f^i :: Dom)
funcType f: D_1 −> D_2	(**define** f^i :: (−> D_1 D_2))
funcType f: Prod(D_1, . . ., D_n) −> D with $n \geq 2$	(**define** f^i :: (−> D_1 . . . D_n D))

Table 2: T_f: Mapping schema of the ASM function declarations to Yices at state i

ASM term	**Yices**
Boolean term: b with $b \in$ {true, false}	b
Integer term: h with $h \in \mathbb{Z}$	h
Natural term: hn with $h \in \mathbb{N}$	h
Enumeration term: E	E
Location term: f	f^i
Location term: $f(a_1, \ldots, a_n)$ with $n \geq 1$	(f^i $\mathsf{T}_t(a_1, i)$ . . . $\mathsf{T}_t(a_n, i)$)
if guard **then** Tthen **else** Telse **endif**	(**if** T_t(guard, i) T_t(Tthen, i) T_t(Telse, i))
(**forall** x_1 **in** D_1, . . ., x_n **in** D_n **with** cond[x_1, . . ., x_n])	(**and** c_1 . . . c_m) with $m = \prod_{j=1}^{n} \lvert D_j \rvert$ where for each $\overline{d}_k = (d_1, \ldots, d_n) \in D_1 \times \ldots \times D_n$: $c_k = \mathsf{T}_t(\text{cond}[x_1 \mapsto d_1, \ldots, x_n \mapsto d_n], i)$
(**exists** x_1 **in** D_1, . . ., x_n **in** D_n **with** cond[x_1, . . ., x_n])	(**or** c_1 . . . c_m) with $m = \prod_{j=1}^{n} \lvert D_j \rvert$ where for each $\overline{d}_k = (d_1, \ldots, d_n) \in D_1 \times \ldots \times D_n$: $c_k = \mathsf{T}_t(\text{cond}[x_1 \mapsto d_1, \ldots, x_n \mapsto d_n], i)$

Table 3: T_t: Mapping schema of ASM terms to Yices at state i

of translation of a simple ASM model. The tool implementing the translator is available online[2].

Signature. The signature is given by the domains and the function declarations.

The mapping function T_{dom} from ASM domains to Yices types is reported in Table 1. For each ASM enumerative domain, we define a Yices scalar type. Basic type-domains Boolean, Integer, and Natural have a straightforward mapping to a corresponding type in Yices. For each domain, we also provide a constant representing the *undef* value; for each enumerative domain D we add the constant D_undef, while for the Integer an Natural domains we select as undef a value of the domain that cannot be assumed by any function of the model[3]. We are not able to provide an undef value for the Boolean domain and, therefore, we require boolean functions to not assume the undef value. An alternative solution could be to model the Boolean domain as a three-valued enumerative domain, but this would greatly complicate the translation. Moreover, since in ASMs rule guards are *formulas* [14] (note 51 at page 64) that need to be always defined, boolean functions used in rule guards cannot be undef: therefore, the limitation of the mapping is limited to boolean functions not used in guards.

Table 2 reports the mapping T_f from ASM function declarations to the corresponding Yices definitions in terms of uninterpreted functions. Note that ASM function declarations must be translated for a generic state i, since they could be added to the logical context multiple times for different states: for this reason, T_f is parametrized with the state i that is reported in the mapped function name. Both 0-ary functions and n-ary functions (with $n > 0$) have a straightforward representation in terms of Yices definitions. Note that there is no difference in the mapping of the declaration of different types of ASM functions (i.e., controlled, monitored, derived, static); only the way to determine their values is mapped in different ways. Static and derived functions have their own definition mapped as Yices definitions (see next but one section), while updates of controlled functions are mapped in the translation of transition rules. In ASMs, monitored functions are not updated by transition rules, but their value is determined by the environment; therefore, their mapping does not require anything but the declaration: in this way, they can assume any value of their domain. This allows us to easily model the *external* nondeterminism due to the environment.

Terms. Both function definitions and transition rules (whose mapping is described in the following two sections) contain terms. Table 3 shows the mapping T_t that maps ASM terms in Yices. Since terms can contain function names, T_t is parametrized with state i.

The mapping of Boolean, Natural, Integer, and enumeration terms is straightforward. An ASM location term is mapped as a Yices function application, where function arguments are the translation of the location parameters values. The conditional term (7th row in Table 3) is mapped in a conditional expression. Forall and exists terms are respectively mapped as conjunction and disjunction of the condition *cond* instantiated over all the possible tuples $d_1, \ldots, d_n$. Although in ASMs the two terms can quantify over infinite domains, the AsmetaL language requires domains $D_1, \ldots, D_n$ to be finite: therefore, the described mapping in SMT is feasible, although the obtained formula could be particularly big in case the domains are big[4]. Note that we could have mapped these terms using the Yices `forall` and `exists` quantified expressions; however, since Yices is not complete when quantifiers are used[5], a Yices context

[2]The tool can be downloaded from http://asmeta.sourceforge.net/download/asm2SMT.html.

[3]Note that, in general, we cannot statically determine such value.

[4]However, since usually the translation in SMT is done for verification purposes, in order to keep the execution time reasonable the domains are usually not too big.

[5]http://yices.csl.sri.com/old/language.shtml#language_quantifiers

containing quantifiers often cannot be evaluated by the solver (it returns the unknown result). Therefore, we preferred to adopt this more verbose mapping that, however, can be always evaluated.

Derived and static functions. In ASMs, besides dynamic functions, the user can introduce derived functions defined by a law over the current state and static functions defined by a law over a number of parameters. These functions cannot be updated by transition rules, but they come with their own definition (i.e., they are functions in the mathematical sense). Their definition is mapped by $\mathsf{T_d}$ as shown in Table 4. For a 0-ary function, the mapping simply consists in creating equality of the function name with the translation of the function body (using the mapping function of terms $\mathsf{T_t}$ shown in Table 3). For n-ary functions (with $n > 0$), instead, the mapping consists in the conjunction of the mapping of the single locations. Note that, also in this case, we could have mapped these function definitions using the Yices `forall` quantified expression; however, as seen before, we could risk to obtain a formula that cannot be always evaluated.

Transition rules. Table 5 reports the mapping function $\mathsf{T_r}$ for transitions rules. It produces a formula that symbolically represents the ASM transition relation. The formula is built by recursively applying the mapping starting from the main rule (last row in the table); the formula is usually **assert**ed in the Yices context in order to represent a generic ASM step (more details on this will be given in Sect. 3.2).

In an *update rule* (first row of Table 5), the location term on the left-hand side of the rule refers to the *next* state, while the term on the right-hand side of the rule refers to *current* state. For this reason, the updated term is mapped with parameter $i + 1$, whereas the term on the right side is mapped with parameter i.

An ASM *parallel rule* for the parallel execution of a set of rules, is mapped as conjunction of the mappings of the single rules.

A complete ASM *conditional rule* (with the else branch) is mapped using the Yices conditional expression, while a partial conditional rule (without the else branch) is mapped as an implication.

An ASM *forall rule* (5th row in Table 5) requires that rule R is executed in parallel with all the values of $x_1 \dots x_n$ that make *guard* true. The rule is mapped as a conjunction of implications stating that if *guard* evaluated with values $d_1 \dots d_n$ holds, then also the mapping of rule R (instantiated with values $d_1 \dots d_n$) must hold. As said before for the forall term and for function definitions, we could use the Yices `forall` quantifier for the translation: however, the obtained Yices context would often produce the unknown result when evaluated.

An ASM *choose rule* requires that rule R is executed once with some values of $x_1 \dots x_n$ (nondeterministically chosen) that make *guard* true (if any). The mapping (6th row in Table 5) first creates a definition cv_j^i for each logical variable x_j of the choose rule[6]; then, it expresses the choose rule through an implication stating that if a tuple of values $d_1 \dots d_n$ exists that make the *guard* true, then the translation of the *guard* and of R must hold. Note that the nondeterminism of the choose rule semantics (*internal* nondeterminism) is compactly embedded in the Yices symbolic representation: any model of the Yices formula (in terms of $cv_1^i \dots cv_n^i$) is a tuple $d_1 \dots d_n$ that allows to execute R[7]: we exploit this feature in the runtime verification of nondeterministic systems described in Sect. 5.

An alternative version of the choose rule allows to specify a rule *Ro* that must be executed if R cannot be executed because *guard* cannot be true. The mapping in Yices (7th row in Table 5) is similar to the previous one, but a conditional expression is used instead of the implication: the mapping of rule *Ro* is specified as else expression.

Similarly to what happens in the translation of the forall and exists terms, the size of the formula obtained by the translation of the forall and choose rules grows with the size of domain $D_1, \dots, D_n$; however, the translation is always feasible since the domains are required to be finite by the starting notation AsmetaL.

3.1 Computation step semantics

In ASMs, a computation step is performed by evaluating the transition rules (starting from the main rule), collecting (in the *update set*) all the updates of controlled locations, and applying the updates. Controlled locations that are not updated keep their value unchanged. The formula obtained by applying the mapping function $\mathsf{T_r}$ to the transition rules (Table 5) correctly determines the update set, but does not guarantee this latter condition. Therefore, for each controlled location f^i, we add the following formula

$$unchLoc_f^i = (\Rightarrow (\textbf{not } (\textbf{or } \mathsf{guard}_1 \dots \mathsf{guard}_n)) (= \mathsf{f}^{i+1} \mathsf{f}^i))$$

being $\mathsf{guard}_1, \dots, \mathsf{guard}_n$ the conditions upon which f^i is updated, and f^{i+1} the location in the next state. Conditions $\mathsf{guard}_1, \dots, \mathsf{guard}_n$ are statically derived from the transition rules that lead to the updates of the location. Let CF_i be all the controlled locations of the ASM model at level i; we define the following formula that asserts the condition for all the locations in CF_i:

$$unchLocs_i = \bigwedge_{f \in CF_i} unchLoc_f^i$$

3.2 Initial state and generic step

We here show how the initial state and a generic step can be represented by using the mapping described in the previous sections.

Let us consider an ASM $M = \langle sig, funcDefs, funcInit, r\_main \rangle$, where *sig* is the signature, $funcInit = \{fi_1, \dots, fi_p\}$ the function initializations and $funcDefs = \{fd_1, \dots, fd_q\}$ the function definitions. We define predicates *init* and $step_i$ to formalize the initial state and the generic step i of the machine, as follows:

$$\text{init} = (\textbf{and } \mathsf{T_d}(\mathsf{fi}_1) \dots \mathsf{T_d}(\mathsf{fi}_p))$$
$$\text{step}_i = (\textbf{and } \mathsf{T_r}(\mathsf{r\_main}, i)\ unchLocs_i\ \mathsf{T_d}(\mathsf{fd}_1^i) \dots \mathsf{T_d}(\mathsf{fd}_q^i))$$

The initial state is determined by the mapping of the controlled function initializations. A generic step is determined by the mapping of the main rule, by the condition on the controlled locations that do not have to change their value, and by the definition of the derived functions at state i[8].

[6] Note that also the definitions for choose rules variables must be parametrized by the state i, since they must be created for each step asserted in the logical context.

[7] Actually, in order to fully describe the semantics of the ASM choose rule, we need to avoid that the mapping of R holds when the *guard* cannot be true. This is guaranteed by additional formulas that state that locations that are not updated do not have to change their value, as explained in Sect. 3.1.

[8] Note that transition rules could read the value of derived functions at state i.

ASM function definition	**Yices**
function f = fd	$(= f^i\ \mathsf{T_t}(fd, i))$
function $f(x_1$ **in** $D_1, \ldots, x_n$ **in** $D_n) = fd[x_1, \ldots, x_n]$ with $n \geq 1$ and $D_1=\{d_1^1, \ldots, d_{m_1}^1\} \ldots D_n=\{d_1^n, \ldots, d_{m_n}^n\}$	$(\textbf{and } (= \mathsf{T_t}(f(d_1^1, \ldots, d_1^n), i)\ \mathsf{T_t}(fd[x_1 \mapsto d_1^1, \ldots, x_n \mapsto d_1^n], i)\ )$ $\ldots (= \mathsf{T_t}(f(d_{m_1}^1, \ldots, d_{m_n}^n), i)\ \mathsf{T_t}(fd[x_1 \mapsto d_{m_1}^1, \ldots, x_n \mapsto d_{m_n}^n], i)\ )\ )$

Table 4: $\mathsf{T_d}$: Mapping schema of ASM function definitions to Yices at state i

ASM transition rule	**Yices**
updatedLoc := updTer	$\mathsf{T_t}(updatedLoc, i+1) = \mathsf{T_t}(updTer, i)$
par $R_1 \ldots R_n$ **endpar**	$(\textbf{and } \mathsf{T_r}(R_1, i) \ldots \mathsf{T_r}(R_n, i))$
if guard **then** Rthen **else** Relse **endif**	$(\textbf{if } \mathsf{T_t}(guard, i)\ \mathsf{T_r}(Rthen, i)\ \mathsf{T_r}(Relse, i))$
if guard **then** Rthen **endif**	$(=> \mathsf{T_t}(guard, i)\ \mathsf{T_r}(Rthen, i))$
forall x_1 **in** $D_1, \ldots, x_n$ **in** D_n **with** $guard[x_1, \ldots, x_n]$ **do** $R[x_1, \ldots, x_n]$	$(\textbf{and } r_1 \ldots r_m)$ with $m = \prod_{j=1}^{n} \lvert D_j \rvert$ where for each $\overline{d}_k = (d_1, \ldots, d_n) \in D_1 \times \ldots \times D_n$ $r_k = (=> \mathsf{T_t}(guard[x_1 \mapsto d_1, \ldots, x_n \mapsto d_n], i)$ $\mathsf{T_r}(R[x_1 \mapsto d_1, \ldots, x_n \mapsto d_n], i)\ )$
choose x_1 **in** $D_1, \ldots, x_n$ **in** D_n **with** $guard[x_1, \ldots, x_n]$ **do** $R[x_1, \ldots, x_n]$	for each x_j: $(\textbf{define } cv_j^i :: D_j)$ $(=> \mathsf{T_t}(\textbf{exists } x_1 \text{ in } D_1, \ldots, x_n \text{ in } D_n \textbf{ with } guard[x_1, \ldots, x_n], i)$ $(\textbf{and } \mathsf{T_t}(guard[x_1 \mapsto cv_1^i, \ldots, x_n \mapsto cv_n^i], i)$ $\mathsf{T_r}(R[x_1 \mapsto cv_1^i, \ldots, x_n \mapsto cv_n^i], i)\ )\ )$
choose x_1 **in** $D_1, \ldots, x_n$ **in** D_n **with** $guard[x_1, \ldots, x_n]$ **do** $R[x_1, \ldots, x_n]$ **otherwise** Ro	for each x_j: $(\textbf{define } cv_j^i :: D_j)$ $(\textbf{if } \mathsf{T_t}(\textbf{exists } x_1 \text{ in } D_1, \ldots, x_n \text{ in } D_n \textbf{ with } guard[x_1, \ldots, x_n], i)$ $(\textbf{and } \mathsf{T_t}(guard[x_1 \mapsto cv_1^i, \ldots, x_n \mapsto cv_n^i], i)$ $\mathsf{T_r}(R[x_1 \mapsto cv_1^i, \ldots, x_n \mapsto cv_n^i], i)\ )$ $\mathsf{T_r}(Ro)\ )$
main rule r_main = mainBody	$\mathsf{T_r}(mainBody, i)$

Table 5: $\mathsf{T_r}$: Mapping schema of ASM transition rules to Yices at state i

3.3 Example of SMT translation

As an example, we consider a tank that can be either filled or emptied. At every step, the tank level can be increased/decreased of up to 3 units of product. The tank is full when it contains 50 units of product. Such tank can be modeled by a simple ASM, as shown in Code 1. The function level records the number of units in the tank; in the initial state the tank is empty. Boolean function full signals whether the tank is full; the function is derived because its value depends on the value of function level. In the main rule, a choose rule nondeterministically increments/decrements the level of the tank of at most three units at a time, not exceeding the maximum capacity. Code 2 shows the translation of the ASM code in Yices. The Yices context contains the definition of the initial state and of a step starting from the initial state. Note that such context represents all the possible computations that can be done in one step.

```
asm Tank
import StandardLibrary
signature:
  domain Level subsetof Integer
  domain IncrDom subsetof Integer
  dynamic controlled level: Level
  derived full: Boolean
definitions:
  domain Level = {0..50}
  domain IncrDom = {−3..3}
  function full = (level = 50)
main rule r_Main =
    choose $x in IncrDom with level + $x >= 0 and level + $x <= 50 do
      level := level + $x
default init s0:
  function level = 0
```

Code 1: ASM model of the Tank case study

4 REFINEMENT PROOF

One of the key concepts of the ASM method is *model refinement* that prescribes that a complete system model should be obtained through a chain of refined models: starting from a high-level abstract description of the system, more precise models should be

```
;; signature state 0
(define level0::(subrange 0 50))
(define full0::bool)
;; initial state
(assert (= level0 0))
;; logic variables for choose rule − step 0
(define x0State0::(subrange −3 3))
;; signature state 1
(define level1::(subrange 0 50))
;; step 0
(assert (and
    ;; transition rules
    (=> (or
          (and (>= (+ level0 −3) 0) (<= (+ level0 −3) 50))
          (and (>= (+ level0 −2) 0) (<= (+ level0 −2) 50))
          (and (>= (+ level0 −1) 0) (<= (+ level0 −1) 50))
          (and (>= level0 0) (<= level0 50))
          (and (>= (+ level0 1) 0) (<= (+ level0 1) 50))
          (and (>= (+ level0 2) 0) (<= (+ level0 2) 50))
          (and (>= (+ level0 3) 0) (<= (+ level0 3) 50))
      (and
        (and (>= (+ level0 x0State0) 0) (<= (+ level0 x0State0) 50))
        (= level1 (+ level0 x0State0))))
    ;; unchanged locations
    (=>
      (not (and (>= (+ level0 x0State0) 0) (<= (+ level0 x0State0) 50)))
      (= level0 level1)
    )
    ;; derived functions definitions
    (= full0 (= level0 50))
 ))
```

Code 2: Yices context of the Tank model (one simulation step)

obtained by iteratively adding more details (possibly reaching a model that is very close to the implementation). A notion of correct refinement has been originally presented in [13]; given an abstract model A and a refined model R, a proof of correct refinement must be able to associate each R-run with some A-run, according to some desired schema (1-1, 1-m, or n-m) and a conformance relation between abstract and refined states. Note that refinement proof checks a property related to the *construction* of the model, i.e., that the model has been correctly refined. Such property is independent of the particular model; model-dependent properties regarding the *behaviour* of the model should be specified using temporal logics and verified using the ASMETA model checker [1].

Automatically proving the original notion of refinement would require to explicitly represent ASM runs of the two machines and compare them using, for example, techniques as model checking. However, developing a technique able to prove any refinement schema for any model is almost impossible. Actually, in our experience in modeling with ASMs [6, 8], we observed that a particular type of 1-m refinement occurs. We called it *stuttering refinement*:

Definition 4.1 (Stuttering Refinement). An ASM R is a correct *stuttering refinement* of an ASM A if and only if each R-run can be split in a sequence of subruns $\widetilde{\rho}_0, \widetilde{\rho}_1, \ldots$ and there is an A-run $S_0, S_1, \ldots$ such that for each $\widetilde{\rho}_i$ it holds $\forall \widetilde{S} \in \widetilde{\rho}_i : conf(\widetilde{S}, S_i)$.

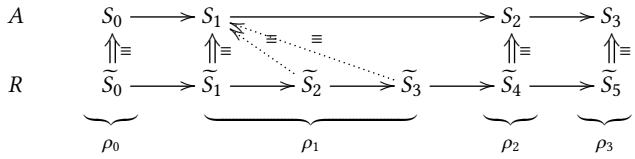

Figure 1: Stuttering refinement: relation between runs

Predicate *conf* formalizes the conformance relation between states of the abstract machine and states of the refined machine:

Definition 4.2 (Conformance). Let S be a state of the abstract machine A (called *abstract state*), $\widetilde{S}$ a state of the refined machine R (called *refined state*). The two states are *conformant* iff corresponding *locations of interest* have equivalent values, i.e.,

$$conf(\widetilde{S}, S) \quad \text{iff} \quad \forall l_R \forall l_A \, corrLoc(l_R, l_A) \rightarrow [\![l_R]\!]_{\widetilde{S}} = [\![l_A]\!]_S$$

where *corrLoc* is a one-to-one correspondence between the locations of interest (i.e., locations on which compare the states) of A and R.

Fig. 1 shows the correspondence between a refined R-run and an abstract A-run.

Proving refinement with SMT. Proving stuttering refinement, differently from the more general notion of refinement [13], does not require to explicitly represent the abstract and refined runs, but can be reduced to the proof of two properties, similarly to what is done in [20] for compiler verification.

Theorem 4.3. *If the following properties hold*

$$\forall \widetilde{S} : (init(\widetilde{S}) \rightarrow \exists S : (init(S) \wedge conf(\widetilde{S}, S))) \tag{1}$$

$$\forall \widetilde{S} \forall \widetilde{S}' \forall S : \left[\begin{pmatrix} step(\widetilde{S}, \widetilde{S}') \\ \wedge \\ conf(\widetilde{S}, S) \end{pmatrix} \rightarrow \begin{pmatrix} \exists S' : (step(S, S') \wedge conf(\widetilde{S}', S')) \\ \vee \\ conf(\widetilde{S}', S) \end{pmatrix} \right] \tag{2}$$

then R is a stuttering refinement of A. $init(S)$ holds iff S is an initial state, and $step(S, S')$ holds iff the state S' can be reached by S in one simulation step.

We name property (1) as *initial refinement*, and property (2) as *step refinement*. Such properties are expressed in terms of states of the abstract and refined machines; we must reduce them to a symbolic representation in order to be able to express them using our encoding. In order to do this, we give the following definitions.

Let $\bar{f}_A = [fa_1, \ldots, fa_n]$ be the functions of the abstract machine A and $\bar{f}'_A = [fa'_1, \ldots, fa'_n]$ their renamed copy in the next state. In the same way, $\bar{f}_R = [fr_1, \ldots, fr_m]$ and $\bar{f}'_R = [fr'_1, \ldots, fr'_m]$ are functions of the refined machine R. In these lists, the first L functions are the locations of interest. We split a list of functions $\bar{f}$ between the functions corresponding to locations of interest and those which are not related in the conformance relation: $\bar{f} = \bar{f}^c + \bar{f}^{nc}$.

Now, we can express the predicates *init*, *step*, and *conf* used in Thm. 4.3, in terms of the function lists of the abstract and refined machines, and rewrite Formulas 1 and 2 as follows[9]:

$$init_R(\bar{f}_R) \rightarrow (\exists \bar{f}^{nc}_A : init_A(\bar{f}^c_R + \bar{f}^{nc}_A)) \tag{3}$$

[9] In [5], we show how Formulas 3 and 4 are derived from Formulas 1 and 2.

```
(define fr_1 :: Dr_1) . . . (define fr_m :: Dr_m)
(define init_R :: bool init_R(f̄_R))
(define existsInit_A :: bool
                (exists ((fa_{L+1} Da_{L+1}) . . . (fa_n Da_n)) init_A(f̄_R^c + f̄_A^nc)) ) )
(assert (not (=> init_R existsInit_A)) ) )
(check)
```

Code 3: Yices context for proving initial refinement

```
(define fr_1 :: Dr_1) . . . (define fr_m :: Dr_m)
(define fr'_1 :: Dr_1) . . . (define fr'_m :: Dr_m)
(define fa_{L+1} :: Da_{L+1}) . . . (define fa_n :: Da_n)
(define step_R :: bool step_R(f̄_R, f̄'_R))
(define existsStep_A :: bool
    (exists ((fa'_{L+1} Da_{L+1}) . . . (fa'_n Da_n)) step_A(f̄_R^c + f̄_A^nc, f̄_R^c' + f̄_A^nc')) ) )
(define stutteringState :: bool (and (= fr'_1 fr_1) . . . (= fr'_L fr_L)) ) )
(assert (not (=> step_R (or existsStep_A stutteringState)) ) ) )
(check)
```

Code 4: Yices context for proving step refinement

$$step_R(\bar{f}_R, \bar{f}'_R) \rightarrow \left(\begin{array}{c} \exists \bar{f}_A^{nc\prime} : step_A(\bar{f}_R^c + \bar{f}_A^{nc}, \bar{f}_R^{c\prime} + \bar{f}_A^{nc\prime}) \\ \vee \\ \bar{f}_R^{c\prime} = \bar{f}_R^c \end{array} \right) \quad (4)$$

The two properties can be mapped in Yices using the mapping described in Sect. 3. Therefore, the automatic proof of stuttering refinement can be reduced to a satisfability checking problem i.e., to two Yices contexts[10]. The two symbolic properties have two useful characteristics w.r.t. the formulation in terms of states:

- the universal quantifications over states are substituted by the introduction of new function symbols (by Herbrandization); such formulation is suitable for checking validity with a logical solver;
- since conformance between abstract and refined states consists in the equality between functions having the same name in the abstract and in the refined machine, we can express the conformance relation by simply using only one copy of the variables. Such formulation is convenient, since it reduces the number of functions in the logical context and also the complexity of the formula, so reducing the complexity of the problem.

For Formula 3, we build the Yices context shown in Code 3. The antecedent of the implication is represented through the SMT definition init_R, and the consequent by definition $\mathsf{existsInit}_A$.

For Formula 4, we build the Yices context shown in Code 4. The antecedent of the implication is represented by definition step_R. The consequent is represented by the disjunction of definitions $\mathsf{existsStep}_A$ and stutteringState; the latter one models the equality of vectors in the stuttering state (i.e., $\bar{f}_R^{c\prime} = \bar{f}_R^c$ in Formula 4) as a conjunction of equalities.

As usual, since we need to prove validity of formulas, we check that their negation is unsatisfiable. Therefore, if the two Yices instances are proved unsatisfiable, the refinement is proved correct. However, step refinement condition is sufficient but not necessary: when Formula 4 is proved not valid (i.e., the corresponding SMT instance is satisfiable), we cannot state that the refinement is not correct. Nevertheless, when refinement is not proved correct, the SMT solver returns a model that we can inspect to understand whether it is really the case that the refinement is not correct, or it is a false negative result (i.e., the model represents a state that is not reachable); in the latter case, we can strengthen the properties using an invariant that restricts the set of states that must be inspected in the proof [5].

[10]The tool implementation of the approach is available at http://asmeta.sourceforge.net/download/asmrefprover.html.

5 RUNTIME VERIFICATION

Another use of the encoding proposed in Sect. 3, is for runtime verification of Java code w.r.t. its ASM specification. This idea, implemented in a tool called CoMA [3], consists in observing the behaviour of a Java object O_C during its execution and checking that its observed elements conform to the expected state of its specification ASM_C. A suitable *link* must be established between O_C fields and pure methods (i.e., methods that cannot change O_C state) and the ASM_C state; this can be done by code annotation. This link identifies the conformance relation (in brief *conf*) among O_C and ASM_C states. Moreover, the designer must designate a set of methods, called *changing methods*, whose invocation corresponds to an ASM step. The runtime verification can be performed by checking that the following schema holds:

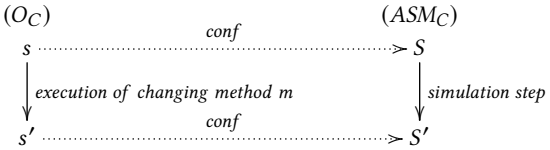

In case of deterministic behaviour (of both the specification and the implementation), runtime verification requires that the next state s' of O_C conforms to the next state S' of the specification. Conformance can be verified by simulation: if O_C conforms to ASM_C in the current state, observed the next state s' of O_C, the monitor checks that it conforms to the next state S' of ASM_C, which can be obtained by applying its transition rules. During simulation, the monitor stores and updates only the current state S of the specification.

In case of nondeterministic behaviours, runtime verification can become complex since the specification takes into account all the possible correct system evolutions. The nondeterminism due to monitored quantities (e.g., the system inputs or external actions), called *external*, is still easy to monitor: once these quantities are fixed by the environment, the system behaves deterministically. However, in most cases, the specification is *internally* nondeterministic, sometimes even when the system is deterministic. For instance, when the system makes random choices, the specification must be nondeterministic. Moreover, nondeterministic specifications can be useful when the modeler wants to be more abstract (with less implementation details) than the corresponding system. Nondeterminism in the behavioural specifications can simplify the representation of complex functionalities and achieve a better separation of concerns [10].

When the behaviour is nondeterministic, the monitor must ensure that the relation shown in Fig. 2 holds. Formally, it can be defined as follows.

Definition 5.1. **Conformant set** Given a Java object O_C, let s_n be the state obtained after n executions of changing methods. We

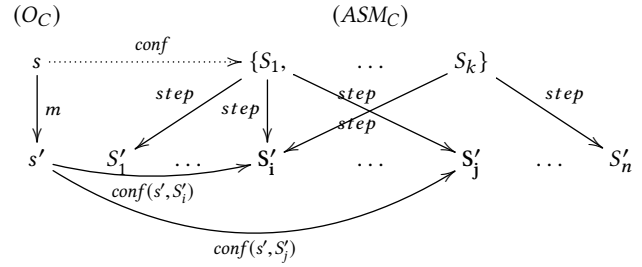

Figure 2: Conformance checking with nondeterminism

call $confSet(s_n)$ the set of ASM states reachable in n steps and conformant with s_n.

Definition 5.2. **Runtime conformance** A class C is runtime conforming to its specification ASM_C if the following conditions hold:

1) the initial state s_0 of the computation of O_C conforms to *at least one* initial state S_0 of the computation of ASM_C, i.e., $\exists S_0$ initial state of ASM_C such that $conf(s_0, S_0)$;
2) for every change step (s, m, s') with s the current state of O_C, for each $S \in confSet(s)$ $\exists (S, S')$ step of ASM_C, such that $conf(s', S')$.

Runtime conformance can be checked as shown in Fig. 2: assuming the current Java state s is conformant with the ASM states $confSet(s) = \{S_1, \ldots, S_k\}$; the Java state s' is produced by the execution of the method m at the state s; ASM states $S'_1, \ldots, S'_n$ are reachable in one step from $confSet(s)$; the Java state s' is checked conformant with at least one state S'_j.

Example 5.3. Let us consider the Tank case study. If the conformance link between the ASM specification and the Java program is only based on the value of function full (since the level value is not observable, for example), then the ASM is nondeterministic. At each step, the ASM has between 4 and 7 possible next states; at most one next state can have value *true* for full. Therefore, if the implementation is correct and the value of full is *false*, more than one of the possible next ASM states can be conformant with the implementation.

Dealing with this kind of monitoring in an explicit state approach would require to keep track of all the possible states to which the monitored system can be conformant. At the ith step of monitoring, the framework should record in the set $confSet(s_i)$ (see Def. 5.1) the ASM states reachable in i steps of simulation that are conformant with the current Java state (as proposed in [16]). If $confSet(s_i)$ becomes empty, then an error is found.

Exploiting the Yices representation of the ASMs, we reduce the runtime conformance checking in the presence of nondeterminism to the satisfiability checking of an STM problem. We symbolically represent the set of states RS_i reachable in i steps of the ASM execution, and the transition relation induced by the ASM transition rules between states in RS_i and their successor states in RS_{i+1}. These formulas establish the logical *context*.

Algorithm 1 CoMA-SMT: monitoring procedure

```
1:  (assert init)                                   ▷ Context initialization
2:  i ← 0
3:  while o_C.m() is invoked do
4:      (assert step_i)                             ▷ Extend context at level i
5:      s_Java ← [|o_C|]                            ▷ Observed Java state after step i
6:      javaValConstr ← getValues(s_Java)           ▷ Get observed values
7:      (assert javaValConstr)                      ▷ Add observed values
8:      if (check) = UNSAT then                     ▷ Is SAT?
9:          return NotConformantException
10:     end if
11:     i ← i + 1
12: end while
```

Example 5.4. Let us consider the Tank case study. The ASM states reachable in one step can be symbolically described as level = 0 ∧ (level − 3 ≤ level′ ≤ level + 3), where level′ represents the updated version of level.

Informally, in order to perform the conformance checking, we build a logical context with the initial state and then we extend the context by asserting a set of formulas stating the values of the observed elements in the implementation current state. Yices is used to check the satisfiability of the obtained context. If the context becomes unsatisfiable, then the implementation is not conformant. Alg. 1 depicts the monitoring procedure of the proposed approach (called CoMA-SMT).

At the beginning, the framework initializes the context (line 1) with the initial state (see Sect. 3.2). The monitoring consists of a never ending loop in which, when a Java changing method m is executed (line 3), the following actions are executed:

- the context is extended for describing the transition relation between ASM states at the current level i and the possible next states at level $i + 1$ (line 4) by $step_i$ as defined in Sect. 3.2;
- from the Java state s_{Java}, obtained after the changing method execution (line 5), and from the linking between the specification and the code, the framework builds the formula *javaValConstr* in which the linked ASM locations are forced to assume the actual values of the corresponding Java elements (line 6). Let $f_1^{i+1}, \ldots, f_g^{i+1}$ be the locations linked to Java fields or methods (i.e., the observed elements) and $v_1, \ldots, v_g$ the values of the linked fields and methods at state $i + 1$. Formula *javaValConstr* is built as follows:

$$(\textbf{and } (= f_1^{i+1}\ v_1) \ldots (= f_g^{i+1}\ v_g))$$

- formula *javaValConstr* is asserted in the logical context (line 7);
- the logical context is checked for satisfiability (line 8):
  - if the context is unsatisfiable, it means that the implementation is not conformant with the specification. In this case, the monitoring is interrupted by throwing an error message (line 9);
  - otherwise, if the context is still satisfiable, it means that the implementation is conformant and the monitoring can continue.

Experiments have shown that CoMA-SMT (i.e., the SMT-based runtime verification approach) has better performances than CoMA (i.e., the explicit state approach) when the degree of nondeterminism is high [4]. However, the introduced overhead is still not negligible

(around 5 ms per step for the Tank case study); moreover, experiments have also shown that CoMA-SMT does not scale well when the logical context grows after each simulation step. Therefore, techniques should be devised in order to keep the size of the logical context (and, therefore, the computation time) reasonable.

6 CONCLUSION AND FUTURE WORK

We have presented a translation of ASMs to SMT that we used for two purposes, namely proving refinement and runtime conformance checking in the presence of nondeterminism. Our experience with SMT is very positive: in the refinement correctness prover, we only needed the formalization of a single step since the proof is inductive, while in the runtime checker we were able to represent (step by step) the reachable conformant states thanks to the support for incremental solving of the SMT solver. The use of the declarative style of SMT theories allowed us to easily deal with nondeterministic behaviours by not having to explicitly keep track of all the possible states. One problem we faced is how to express what does not change (see Sect. 3.1) and, in this, operational notations are still better than declarative approaches. Note that we could use BDDs to symbolically represent ASM states; however, for bounded model checking it has been shown that a SAT/SMT approach scales better [11, 19]. So, since both our runtime verification and refinement approaches have several commonalities with BMC, we believe that the SMT approach has still several advantages.

As future work, we plan to employ our encoding in other V&V activities. In [2], we presented an approach for doing *model review* of ASMs that checks some *meta-properties* that any model should guarantee. The approach is sound and complete but, since based on model checking, does not scale well. As future work, we plan to detect possible inconsistencies of the ASM model using the SMT representation (e.g., an inconsistent context is a signal of the presence of inconsistent updates); that approach could be not always complete: it could only find a meta-property violation, but fail in proving their validity. However, we believe that SMT would scale much better than the pure model checking approach and we plan to apply induction for proving meta-properties.

In [9], we presented a formalization of self-adaptive systems in terms of ASMs. One problem with self-adaptive systems is how to resolve conflicts of different adaptation scenarios at runtime; in general, among n adaptation scenarios that are applicable in a given state, $m \leq n$ scenarios can simultaneously fire without conflicts. We plan to exploit the SMT representation to find the maximum subset of adaptation scenarios that can simultaneously fire.

Finally, we plan to exploit the symbolic representation for checking model-dependent behavioural properties, in a classical formal verification approach like that presented in [21].

REFERENCES

[1] Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2010. AsmetaSMV: A Way to Link High-Level ASM Models to Low-Level NuSMV Specifications. In *Abstract State Machines, Alloy, B and Z (Lecture Notes in Computer Science)*, Marc Frappier, Uwe Glässer, Sarfraz Khurshid, Régine Laleau, and Steve Reeves (Eds.), Vol. 5977. Springer Berlin Heidelberg, 61–74. https://doi.org/10.1007/978-3-642-11811-1_6

[2] Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2010. Automatic Review of Abstract State Machines by Meta Property Verification. In *Proceedings of the Second NASA Formal Methods Symposium (NFM 2010), NASA/CP-2010-216215*, César Muñoz (Ed.). NASA, Langley Research Center, Hampton VA 23681-2199, USA, 4–13.

[3] Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2012. CoMA: Conformance Monitoring of Java Programs by Abstract State Machines. In *Runtime Verification*, Sarfraz Khurshid and Koushik Sen (Eds.). Lecture Notes in Computer Science, Vol. 7186. Springer Berlin Heidelberg, 223–238. https://doi.org/10.1007/978-3-642-29860-8_17

[4] Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2014. Using SMT for dealing with nondeterminism in ASM-based runtime verification. *ECEASST* 70 (2014). https://doi.org/10.14279/tuj.eceasst.70.970

[5] Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2016. SMT-based automatic proof of ASM model refinement. In *Software Engineering and Formal Methods: 14th International Conference, SEFM 2016, Held as Part of STAF 2016, Vienna, Austria, July 4-8, 2016, Proceedings*, Rocco De Nicola and Eva Kühn (Eds.). Springer International Publishing, Cham, 253–269. https://doi.org/10.1007/978-3-319-41591-8_17

[6] Paolo Arcaini, Angelo Gargantini, and Elvinia Riccobene. 2017. Rigorous development process of a safety-critical system: from ASM models to Java code. *International Journal on Software Tools for Technology Transfer* 19, 2 (2017), 247–269. https://doi.org/10.1007/s10009-015-0394-x

[7] Paolo Arcaini, Angelo Gargantini, Elvinia Riccobene, and Patrizia Scandurra. 2011. A model-driven process for engineering a toolset for a formal method. *Software: Practice and Experience* 41 (2011), 155–166. Issue 2.

[8] Paolo Arcaini, Roxana-Maria Holom, and Elvinia Riccobene. 2016. ASM-based formal design of an adaptivity component for a Cloud system. *Formal Aspects of Computing* 28, 4 (2016), 567–595. https://doi.org/10.1007/s00165-016-0371-5

[9] Paolo Arcaini, Elvinia Riccobene, and Patrizia Scandurra. 2017. Formal Design and Verification of Self-Adaptive Systems with Decentralized Control. *ACM Trans. Auton. Adapt. Syst.* 11, 4, Article 25 (Jan. 2017), 35 pages.

[10] Pieter Bekaert and Eric Steegmans. 2001. Non-determinism in conceptual models. In *Proceedings of the Tenth OOPSLA Workshop on Behavioral Semantics*. 24 – 34.

[11] Armin Biere. 2009. Bounded Model Checking. In *Handbook of Satisfiability*, Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (Eds.). Frontiers in Artificial Intelligence and Applications, Vol. 185. IOS Press, 457–481.

[12] Silvia Bonfanti, Marco Carissoni, Angelo Gargantini, and Atif Mashkoor. 2017. Asm2C++: A Tool for Code Generation from Abstract State Machines to Arduino. In *NASA Formal Methods: 9th International Symposium, NFM 2017, Proceedings*, Clark Barrett, Misty Davies, and Temesghen Kahsai (Eds.). Springer, 295–301. https://doi.org/10.1007/978-3-319-57288-8_21

[13] Egon Börger. 2003. The ASM Refinement Method. *Formal Aspects of Computing* 15, 2 (2003), 237–257.

[14] Egon Börger and Robert Stärk. 2003. *Abstract State Machines: A Method for High-Level System Design and Analysis*. Springer Verlag.

[15] Bruno Dutertre and Leonardo de Moura. 2006. *The Yices SMT solver*. Technical Report. SRI Available at http://yices.csl.sri.com/tool-paper.pdf.

[16] Yliès Falcone, Klaus Havelund, and Giles Reger. 2013. A Tutorial on Runtime Verification. In *Engineering Dependable Software Systems*, Manfred Broy, Doron A. Peled, and Georg Kalus (Eds.). NATO Science for Peace and Security Series, D: Information and Communication Security, Vol. 34. IOS Press, 141–175. https://doi.org/10.3233/978-1-61499-207-3-141

[17] Angelo Gargantini and Elvinia Riccobene. 2000. Encoding Abstract State Machines in PVS. In *Abstract State Machines - Theory and Applications: International Workshop, ASM 2000 Monte Verità, Switzerland, March 19–24, 2000 Proceedings (Lecture Notes in Computer Science)*, Yuri Gurevich, Philipp W. Kutter, Martin Odersky, and Lothar Thiele (Eds.), Vol. 1912. Springer Berlin Heidelberg, Berlin, Heidelberg, 303–322. https://doi.org/10.1007/3-540-44518-8_17

[18] Yuri Gurevich. 1995. Specification and Validation Methods. Oxford University Press, Inc., New York, NY, USA, Chapter Evolving Algebras 1993: Lipari Guide, 9–36.

[19] Robert P. Kurshan. 2008. Verification Technology Transfer. In *25 Years of Model Checking*, Orna Grumberg and Helmut Veith (Eds.). Lecture Notes in Computer Science, Vol. 5000. Springer Berlin Heidelberg, 46–64. https://doi.org/10.1007/978-3-540-69850-0_3

[20] Kedar S. Namjoshi, Giacomo Tagliabue, and Lenore D. Zuck. 2013. A Witnessing Compiler: A Proof of Concept. In *Runtime Verification: 4th International Conference, RV 2013, Rennes, France, September 24-27, 2013. Proceedings*, Axel Legay and Saddek Bensalem (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 340–345. https://doi.org/10.1007/978-3-642-40787-1_22

[21] Margus Veanes, Nikolaj Bjørner, Yuri Gurevich, and Wolfram Schulte. 2009. Symbolic Bounded Model Checking of Abstract State Machines. *Int. J. Software and Informatics* 3, 2-3 (2009), 149–170.