Christiano Braga and Narciso Martí-Oliet (Eds.)

Preproceedings of SBMF 2014
17th Brazilian Symposium on Formal Methods

Preface

This volume contains the papers presented at SBMF 2014: the 17th Brazilian Symposium on Formal Methods. The conference was held in Maceió, Brazil, as part of CBSoft 2014, the 5th Brazilian Conference on Software: Theory and Practice.

The conference program included two invited talks, given by David Deharbe (UFRN, Brazil) and Narciso Martí-Oliet (Universidad Complutense de Madrid, Spain).

A total of 14 research papers were presented at the conference: 10 full papers and 4 short papers. They are all included in these preproceedings. The authors of full papers will have a chance to revise their papers once more and incorporate feedback received during the conference. Further revised full papers will be published after the conference as a volume of Lecture Notes in Computer Science, by Springer.

The contributions were selected from 34 submissions that came from 18 different countries: Brazil, Canada, Colombia, Denmark, France, Germany, India, Israel, Italy, Pakistan, Portugal, South Africa, Switzerland, Tunisia, Turkey, Ukraine, UK, and Uruguay.

The processes of submission by the authors, paper review and deliberations of the program committee were all assisted by EasyChair.

We are grateful to the program committee and to the referees for their hard work in evaluating submissions and suggesting improvements. SBMF 2014 was organized by the Universidade Federal de Alagoas (UFAL) and promoted by the Brazilian Computer Society (SBC). We are very thankful to the local organizers of CBSoft 2014, that were coordinated by Marcio Ribeiro, Baldoino Santos Neto and Leandro Dias da Silva, all from UFAL, for their very hard work.

We hope you enjoy reading these proceedings as much as we enjoyed preparing them.

September 2014

Christiano Braga and Narciso Martí-Oliet
Program Chairs
SBMF 2014

Organization

Program Commitee

Aline Andrade (UFBA, Brazil)
Wilkerson Andrade (UFCG, Brazil)
Luis Barbosa (Universidade do Minho, Portugal)
Christiano Braga (UFF, Brazil, co-chair)
Michael Butler (University of Southampton, UK)
Ana Cavalcanti (University of York, UK)
Marcio Cornelio (UFPE, Brazil)
Andrea Corradini (University of Pisa, Italy)
Jim Davies (University of Oxford, UK)
David Deharbe (UFRN, Brazil)
Ewen Denney (RIACS/NASA, USA)
Clare Dixon (University of Liverpool, UK)
Jorge Figueiredo (UFCG, Brazil)
Marcelo Frias (Instituto Tecnologico de Buenos Aires, AR)
Rohit Gheyi (UFCG, Brazil)
Juliano Iyoda (UFPE, Brazil)
Zhiming Liu (Birmingham City University, UK)
Patricia Machado (UFCG, Brazil)
Tiago Massoni (UFCG, Brazil)
Ana Melo (USP, Brazil)
Alvaro Moreira (UFRGS, Brazil)
Anamaria Moreira (UFRN, Brazil)
Carroll Morgan (University of New South Wales, Australia)
Arnaldo Moura (UNICAMP, Brazil)
Leonardo Moura (Microsoft Research, USA)
Narciso Marti-Oliet (Universidad Complutense de Madrid, co-chair)
Alexandre Mota (UFPE, Brazil)
David Naumann (Stevens Institute of Technology, USA)
Daltro Nunes (UFRGS, Brazil)
Jose Oliveira (Universidade do Minho, Portugal)
Marcel Oliveira (UFRN, Brazil)
Peter Olveczky (University of Oslo, Norway)
Alberto Pardo (Universidad de la Republica, Uruguay)
Alexandre Petrenko (CRIM, Canada)
Leila Ribeiro (UFRGS, Brazil)
Augusto Sampaio (UFPE, Brazil)
Leila Silva (UFS, Brazil)
Adenilso Simao (ICMC-USP, Brazil)
Heike Wehrheim (University of Paderborn, Germany)
Jim Woodcock (University of York, UK)

Referees

Paulo Salem
Omer Landry Nguena Timo
Simone Hanazumi
Sarah Loos
Arnaud Dury
Vitaly Savicks
Jun Pang
Regina Motz
Alan Moraes
Regivan Santiago
Edward Hermann Haeusler
Andrea Vandin
Daniel Fridlender
Asieh Salehi Fathabadi
Jun Pang

Local Organization

Marcio Ribeiro, Baldoino Santos Neto e Leandro Dias da Silva (UFAL)

Promoting and Sponsoring Institutions

Sociedade Brasileira de Computação
CAPES, CNPq, Ines, and Google

Table of Contents

b2llvm: B developments onto the LLVM

David Deharbe

Universidade Federal do Rio Grande do Norte
david@dimap.ufrn.br

Abstract. We present BLLVM, a multi-platform code generator for the B-method. BLLVM currently handles the following elements of the B language: simple data types, imperative instructions and component compositions. In particular, this paper describes the translation from some essential constructs of the B language for implementations towards LLVM source code. We use an example-based approach for this description.

Equational abstractions in Rewriting Logic and Maude

Narciso Martí-Oliet

Universidad Complutense de Madrid
narciso@ucm.es

Abstract. Maude is a high-level language and high-performance system supporting both equational and rewriting computation for a wide range of applications. Maude also provides a model checker for linear temporal logic. This procedure can be used to prove properties when the set of states reachable from an initial state in a system is finite; when this is not the case, it may be possible to use an equational abstraction technique for reducing the size of the state space. Abstraction reduces the problem of whether an infinite state system satisfies a temporal logic property to model checking that property on a finite state abstract version. The most common abstractions are quotients of the original system. We present a simple method of defining quotient abstractions by means of equations collapsing the set of states. Our method yields the minimal quotient system together with a set of proof obligations that guarantee its executability and can be discharged with tools such as those available in the Maude formal environment. The proposed method will be illustrated in a couple of detailed examples.

Parameterisation of Three-Valued Abstractions

Nils Timm and Stefan Gruner

Department of Computer Science, University of Pretoria, South Africa
{ntimm,sgruner}@cs.up.ac.za

Abstract. Three-valued abstraction is an established technique in software model checking. It proceeds by generating a state space model over the values *true*, *false* and *unknown*, where the latter value is used to represent the loss of information due to abstraction. Temporal logic properties can then be evaluated on such models. In case of an *unknown* result, the abstraction is iteratively refined. In this paper, we introduce *parameterised three-valued model checking*. In our new type of models, unknown parts can be either associated with the constant value *unknown* or with expressions over boolean parameters. Our parameterisation is an alternative way to state that the truth value of certain predicates or transitions is actually not known and that the checked property has to yield the same result under each possible parameter instantiation. A novel feature of our approach is that it allows for establishing logical connections between parameters: While *unknown* parts in pure three-valued models are never related to each other, our parameterisation approach enables to represent facts like 'a certain pair of transitions has unknown but complementary truth values', or 'the value of a predicate is unknown but remains constant along all states of a certain path'. We demonstrate that such facts can be automatically derived from the system to be verified and that covering these facts in an abstract model can be crucial for the success and efficiency of checking temporal logic properties. Moreover, we introduce an automatic verification framework based on counterexample-guided abstraction refinement and parameterisation.

1 Introduction

Predicate abstraction [2] is an established technique for reducing the complexity of temporal logic model checking. It proceeds by generating a state space model of the software system to be analysed. In this model, concrete states of the system are mapped to abstract states over a finite set of predicates, and admissible executions of the system are represented by sequences of transitions between states. Traditional predicate abstraction techniques are based on a boolean domain for predicates and on an over-approximation of the concrete state space. Thus, only universal properties are preserved under this form of abstraction. If checking a universal property for an abstract model yields *false*, it cannot be concluded that the original system violates this property as well. In this case, model checking additionally returns an *abstract counterexample* - a path in the model that refutes the property. In order to gain certainty about whether this counterexample is spurious or corresponds to a real path, it has to be simulated on

the original system. The simulation of counterexamples involves a partial exploration of the concrete state space, and thus, can be exceedingly costly. Spurious counterexamples are typically ruled out via *counterexample-guided abstraction refinement* (CEGAR) [4]: Further predicates over the variables of the system are iteratively added to the model until a level of abstraction is reached where the property can be either definitely proved or a real counterexample can be found. The application of CEGAR does, however, not guarantee that eventually a model can be constructed that is both precise enough for a definite outcome and small enough to be manageable with the available computational resources.

More recent approaches [3, 18, 13] to abstraction refinement for model checking are based on a domain for predicates with the truth values *true, false* and *unknown*. Corresponding three-valued models with the additional value *unknown* enable to explicitly model the loss of information due to abstraction. In comparison to boolean abstractions, the three-valued approach is capable of preserving universal *and* existential properties. Hence, all definite results in three-valued model checking can be directly transferred to the original system. Only an *unknown* result necessitates iterative refinement. In the latter case, an *unconfirmed counterexample* – a potential error path in the model with *unknown* transitions and predicates – is returned. Unconfirmed counterexamples directly hint at necessary refinement steps. Thus, the costly simulation of counterexamples on the original system is not required in the three-valued setting. Model checking three-valued abstractions can be conducted at the same cost as checking boolean abstractions, but it additionally comes along with the aforementioned advantages.

Continuative work in this field has shown that the precision of model checking three-valued abstractions can be increased by the concept of *generalised model checking* (GMC) [7]. While standard three-valued model checking (3MC) [3, 18, 13] is based on a special *three-valued* semantics that enables the direct evaluation of temporal logic formulae on three-valued models, the idea of GMC is to construct *all* boolean concretisations of a three-valued model. Then classical two-valued model checking is applied to each concretisation and it is checked whether the results are consistent, i.e. whether either all results are *true* or whether all are *false*. In case of consistency, the result can be transferred to the original system. GMC generally yields more definite results than 3MC. Hence, the application of GMC instead of 3MC can reduce the number of necessary refinement iterations in abstraction-based verification. However, the 3MC problem is PSPACE-complete, whereas the GMC problem is even EXP-complete: Number and size of concretisations can be exponential in the size of the three-valued model. Thus, GMC is rather of theoretical than of practical interest. Most existing three-valued abstraction-based verification frameworks, e.g. [13, 8, 14], rely on standard 3MC and try to compensate the lack of precision with additional refinement steps.

Here, we introduce *parameterised three-valued model checking* (PMC) which is a hybrid of three-valued and generalised model checking. Predicates and transitions in our parameterised three-valued models can be either associated with the values *true*, *false* or *unknown* – or with expressions over boolean parame-

ters. Our parameterisation is an alternative way to state that the truth value of certain predicates or transitions is actually not known and that the checked property has to yield the same result under each parameter instantiation. PMC is thus conducted via evaluating a temporal logic formula under all parameter instantiations and checking whether the results are consistent. In contrast to GMC, parameterised three-valued model checking reduces to multiple instances of standard three-valued model checking, since the instantiation only affects parameters but not the explicit truth value *unknown*. Sizes of instantiations are always linear in the size of the parameterised three-valued model. Moreover, parameterisation particularly allows to establish logical connections between *unknown*s in the abstract model: While *unknown* parts in 3MC and GMC are never related to each other, our parameterisation approach enables to represent facts like 'a certain pair of transitions has unknown but complementary truth values', or 'the value of a predicate is unknown but remains constant along all states of a certain path'. We demonstrate that such facts can be automatically derived from the software system to be verified and that covering these facts in an abstract model can be crucial for the success and efficiency of checking temporal logic properties. In particular, we introduce an automatic verification framework for concurrent systems based on parameterised three-valued model checking: Starting with pure three-valued abstraction, in each iteration either classical refinement or parameterisation of *unknown* parts is applied until a definite result in verification can be obtained. The decisions for refinement or parameterisation are automatically made based on unconfirmed counterexamples. For several verification tasks our hybrid approach can significantly outperform the pure three-valued approach. Our work includes the definition of parameterisation rules for three-valued abstractions and a proven theorem which states that PMC is sound if parameterisation is applied according to the rules.

2 Background: Three-Valued Model Checking

We start with a brief introduction to three-valued state space models, here three-valued Kripke structures, and the evaluation of temporal logic properties on them. The key feature of these Kripke structures is a third truth value $\bot$ (i.e. *unknown*) for transitions and labellings, which can be used to model uncertainty.

Definition 1 (Three-Valued Kripke Structure). *A* three-valued Kripke structure *over a set of atomic predicates AP is a tuple $K = (S, R, L, \mathbb{F})$ where*

- *S is a finite set of states,*
- *$R : S \times S \to \{true, \bot, false\}$ is a transition function with $\forall\, s \in S : \exists\, s' \in S : R(s, s') \in \{true, \bot\}$,*
- *$L : S \times AP \to \{true, \bot, false\}$ is a labelling function that associates a truth value with each predicate in each state,*
- *$\mathbb{F} \subseteq \mathcal{P}(R^{-1}(\{true, \bot\}))$ is a set of fairness constraints where each constraint $F \in \mathbb{F}$ is a set of* non-false *transitions.*

An example for a Kripke structure K over a set $AP = \{p\}$ is depicted below.

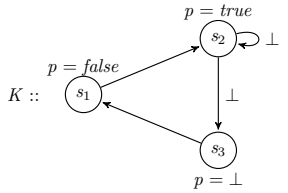

A path π of a three-valued Kripke structure K is an infinite sequence of states $s_1 s_2 s_3 \ldots$ with $R(s_i, s_{i+1}) \in \{true, \bot\}$. π_i denotes the i-th state of π, whereas π^i denotes the i-th suffix $\pi_i \pi_{i+1} \pi_{i+2} \ldots$ of π. A path π is fair if it takes infinitely often a transition from every fairness constraint $F \in \mathbb{F}$. By $\Pi(K, s)$ we denote the set of all fair paths of K starting in $s \in S$. Paths are considered for the evaluation of temporal logic properties of Kripke structures. Here we use the linear temporal logic (LTL) for specifying properties.

Definition 2 (Syntax of LTL). *Let AP be a set of atomic predicates and $p \in AP$. The* syntax of LTL formulae *ψ is given by*

$$\psi ::= p \mid \neg\psi \mid \psi \vee \psi \mid \psi \wedge \psi \mid \mathbf{X}\psi \mid \mathbf{F}\psi \mid \mathbf{G}\psi \mid \psi \mathbf{U} \psi.$$

Due to the extended domain for truth values in three-valued Kripke structures, the evaluation of LTL formulae is not based on classical two-valued logic. In three-valued model checking we operate under the three-valued Kleene logic $\mathbb{K}_3$ [6] whose semantics is given by the truth tables below.

$\wedge$	*true*	$\bot$	*false*
true	*true*	$\bot$	*false*
$\bot$	$\bot$	$\bot$	*false*
false	*false*	*false*	*false*

$\vee$	*true*	$\bot$	*false*
true	*true*	*true*	*true*
$\bot$	*true*	$\bot$	$\bot$
false	*true*	$\bot$	*false*

$\neg$	
true	*false*
$\bot$	$\bot$
false	*true*

For $\mathbb{K}_3$ we have a reflexive *information ordering* $\leq_{\mathbb{K}_3}$ (in words: 'less or equal definite than') with $\bot \leq_{\mathbb{K}_3}$ *true*, $\bot \leq_{\mathbb{K}_3}$ *false*, and *true*, *false* incomparable. Based on $\mathbb{K}_3$, linear temporal logic formulae can be evaluated on paths of three-valued Kripke structures according to the following definition.

Definition 3 (Three-Valued Evaluation of LTL). *Let $K = (S, R, L, \mathbb{F})$ over AP be a three-valued Kripke structure. Then the* evaluation *of an LTL formula ψ on a fair path π of K, written $[\pi \models \psi]$, is inductively defined as follows*

$$\begin{aligned}
[\pi \models p] &:= L(\pi_1, p) \\
[\pi \models \neg\psi] &:= \neg\,[\pi \models \psi] \\
[\pi \models \psi \vee \psi'] &:= [\pi \models \psi] \ \vee \ [\pi \models \psi'] \\
[\pi \models \mathbf{X}\psi] &:= R(\pi_1, \pi_2) \wedge \left[\pi^2 \models \psi\right] \\
[\pi \models \mathbf{G}\psi] &:= \textstyle\bigwedge_{i \in \mathbb{N}} \left(R(\pi_i, \pi_{i+1}) \wedge \left[\pi^i \models \psi\right]\right) \\
[\pi \models \mathbf{F}\psi] &:= \textstyle\bigvee_{i \in \mathbb{N}} \left(\left[\pi^i \models \psi\right] \wedge \bigwedge_{0 \leq j < i} R(\pi_i, \pi_{i+1})\right) \\
[\pi \models \psi \mathbf{U} \psi'] &:= \textstyle\bigvee_{i \in \mathbb{N}} \left(\left[\pi^i \models \psi'\right] \wedge \bigwedge_{0 \leq j < i} \left(R(\pi_j, \pi_{j+1}) \wedge [\pi^j \models \psi]\right)\right)
\end{aligned}$$

The evaluation of LTL formulae on entire three-valued Kripke structures is what we call *three-valued model checking* [3].

Definition 4 (Three-Valued LTL Model Checking). *Let $K = (S, R, L, \mathbb{F})$ over AP be a three-valued Kripke structure. Moreover, let ψ be an LTL formula over AP. The value of ψ in a state s of K, written $[K, s \models \psi]$, is defined as*

$$[K, s \models \psi] \;:=\; \bigwedge_{\pi \in \Pi(K,s)} [\pi \models \psi]$$

In three-valued model checking there exist three possible outcomes: *true*, *false* and $\bot$. Three-valued model checking reduces to classical two-valued model checking if the Kripke structure K is actually two-valued, i.e. $R^{-1}(\bot) = \varnothing$ and $L^{-1}(\bot) = \varnothing$. In this case, only the outcomes *true* and *false* are possible. For our example Kripke structure $[K, s_1 \models \mathbf{G}p]$ yields *false*, whereas $[K, s_1 \models \mathbf{GF}p]$ yields *unknown*. $\mathbf{G}p$ is a temporal logic formula that characterises a typical *safety* property, while $\mathbf{GF}p$ characterises a *liveness* property. Safety and liveness are the most vital requirements in software verification. In our approach, we therefore particularly focus on these two kinds of properties.

For the sake of completeness, we also briefly review generalised model checking (for more details see [7]). Under GMC, $[K, s \models \psi]$ yields *true* iff $[K', s \models \psi]$ is *true* for all concretisations K' of K, where a concretisation is a two-valued K' such that $[K, s \models \psi] \;\leq_{\mathbb{K}_3}\; [K', s \models \psi]$ for all LTL formulae ψ. The definition of $[K, s \models \psi] = \mathit{false}$ is analogous. In all remaining cases $[K, s \models \psi]$ yields $\bot$.

3 Parameterised Three-Valued Model Checking

State space models constructed by three-valued abstraction techniques [13, 8, 14] are typically represented as (pure) three-valued Kripke structures. Here we introduce a generalisation called *parameterised three-valued Kripke structures*, and we define model checking for these structures. Later we will see that *parameterised three-valued model checking* (PMC) for three-valued abstractions can significantly enhance the precision of verification.

Definition 5 (Parameterised Three-Valued Kripke Structure). *A* parameterised three-valued Kripke structure *over AP and a set of boolean parameters $X = \{x_1, \ldots, x_m\}$ is a parameterised tuple $K(\overset{m}{x}) = (S, R(\overset{m}{x}), L(\overset{m}{x}), \mathbb{F}(\overset{m}{x}))$ where*

- *S is a finite set of states,*
- *$R(\overset{m}{x}) : S \times S \to \{\mathit{true}, \bot, \mathit{false}\} \cup BE(X)$ is a transition function with $\forall\, s \in S : \exists\, s' \in S : R(\overset{m}{x})(s, s') \in \{\mathit{true}, \bot\} \cup BE(X)$ where $BE(X)$ denotes the set of boolean expressions over X,*
- *$L(\overset{m}{x}) : S \times AP \to \{\mathit{true}, \bot, \mathit{false}\} \cup BE(X)$ is a labelling function that associates a truth value or a parameter expression with each predicate in each state,*
- *$\mathbb{F}(\overset{m}{x}) \subseteq \mathcal{P}(R^{-1}(\overset{m}{x})(\{\mathit{true}, \bot\} \cup BE(X)))$ is a set of fairness constraints where each constraint $F \in \mathbb{F}(\overset{m}{x})$ is a set of* non-false *transitions.*

Note that $(\overset{m}{x})$ is an abbreviation for the parameter tuple $(x_1, \ldots, x_m)$. An instantiation of a parameterised three-valued Kripke structure $K(\overset{m}{x})$ is a *pure* three-valued Kripke structure $K(\overset{m}{a})$ where $(\overset{m}{a}) \in \{true, false\}^m$. Hence, all parameters are substituted by *boolean* truth values. However, predicates and transitions that were not parameterised in $K(\overset{m}{x})$ may still hold the value *unknown* in $K(\overset{m}{a})$. If the current tuple of parameters or truth values is clear from the context, we will not explicitly mention it, i.e. we will just refer to R, L and $\mathbb{F}$. An example for a parameterised three-valued Kripke structure together with all its pure three-valued instantiations is shown in the figure below.

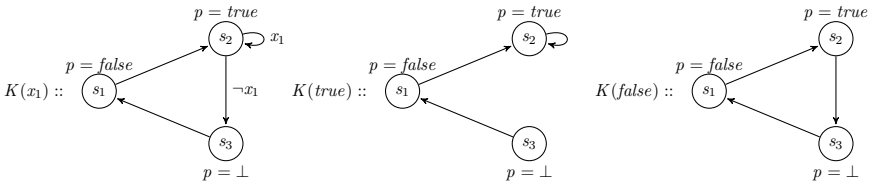

For evaluating temporal logic formulae on parameterised three-valued Kripke structures we consider all possible instantiations.

Definition 6 (Parameterised Three-Valued LTL Model Checking). *Let $K(\overset{m}{x}) = (S, R(\overset{m}{x}), L(\overset{m}{x}), \mathbb{F}(\overset{m}{x}))$ be a parameterised three-valued Kripke structure over AP and $X = \{x_1, \ldots, x_m\}$. Moreover, let ψ be an LTL formula over AP. The value of ψ in a state s of $K(\overset{m}{x})$, written $[K(\overset{m}{x}), s \models \psi]$, is defined as*

$$\left[K(\overset{m}{x}), s \models \psi\right] := \begin{cases} true & \text{if } \bigwedge_{(\overset{m}{a}) \in \{t,f\}^m} \left(\left[K(\overset{m}{a}), s \models \psi\right] = true\right) \\ false & \text{if } \bigwedge_{(\overset{m}{a}) \in \{t,f\}^m} \left(\left[K(\overset{m}{a}), s \models \psi\right] = false\right) \\ \bot & else \end{cases}$$

Thus, if checking a temporal logic property yields *true* for all instantiations, the result is transferred to the parameterised Kripke structure. The same holds for *false* results for all instantiations. In all other cases PMC returns *unknown*. For our recent example, we get $[K(x_1), s_1 \models \mathbf{GF}p] = true$ since $\mathbf{GF}p$ holds for both $K(true)$ and $K(false)$. In contrast to our example from Section 2, the two outgoing transitions of state s_2 are no longer *unknown* but parameterised. Moreover, we capture the fact that the associated transition values are *complementary*, which gives us the necessary precision for a definite result in verification.

Subsequently, we will see that such facts can be automatically derived from the control flow and program code of the modelled system in the sense that the corresponding parameterisation gives us a sound abstraction. Furthermore, we will show how parameterised three-valued model checking can be effectively integrated into an automatic abstraction refinement-based verification procedure.

4 Application to Three-Valued Abstractions

Three-valued model checking [3] is used in many abstraction-based verification frameworks for software systems [13, 10, 8, 1]. An effective state space reduction

technique for concurrent software systems is *three-valued spotlight abstraction* [12, 14, 15]. In previous works [16, 17], we have demonstrated that verifying concurrent systems via spotlight abstraction and three-valued model checking can significantly outperform approaches based on boolean predicate abstraction [2]. In this section, we give a brief introduction to concurrent systems and spotlight abstraction (for more details see [12]). Moreover, we show how *parameterisation* can be applied to three-valued Kripke structures constructed by spotlight abstraction and how this can increase the efficiency of verification.

4.1 Spotlight Abstraction for Concurrent Systems

A concurrent system Sys consists of a number of asynchronous processes composed in parallel: $Sys = \|_{i=1}^{n} Proc_i$. It is defined over a set of variables $Var = Var_s \cup \bigcup_{i=1}^{n} Var_i$ where Var_s is a set of shared variables and $Var_1, \ldots, Var_n$ are sets of local variables associated with the processes $Proc_1, \ldots, Proc_n$, respectively. A process corresponds to a finite sequence of locations where each location is associated with an operation op on the variables in $Var_s \cup Var_i$. Operations are of the form $op = assume(e) : v_1 := e_1, \ldots, v_k := e_k$ where $e, e_1, \ldots, e_k$ are expressions over $Var_s \cup Var_i = \{v_1, \ldots, v_k\}$. Hence, an operation consists of an assume part, also called *guard*, and a list of assignments. Executing the guard blocks the execution of the assignments until the expression e evaluates to *true*. We omit the guard if e is constantly *true*. The current location of a process $Proc_i$ can be regarded as the value of an additional local counter variable pc_i over the process' locations $Loc_i = \{1_i, \ldots, L_i\}$. Locations may also be associated with compound operations, which consist of one or more sub-operations nested inside a control structure. Compound operations in our systems are, amongst others, *if-then-else* and *while-do*. An example for a concurrent system is depicted below.

$$v_1, \ldots, v_k \ : \textbf{integer}$$

$$Proc_1 :: \begin{bmatrix} 1: [\ \ldots\ ] \\ 2: \textbf{while}\ (v_1 > 0)\ \textbf{do} \\ \quad 3: [\ \ldots\ ] \\ 4: \ progress \\ 5: [\ \ldots\ ] \end{bmatrix} \ \| \ Proc_2 :: \begin{bmatrix} 1: [\ \ldots\ ] \\ 2: \ v_1 := f(v_2, \ldots, v_k) \\ 3: [\ \ldots\ ] \end{bmatrix} \ \| \ \ldots \ \| \ Proc_n$$

Here we have a composition of n processes operating on the shared variables $v_1, \ldots, v_k$. A liveness property to verify might be whether $Proc_1$ always repeatedly reaches *progress*, which we assume is an arbitrary assertion over $Proc_1$'s variables. Subsequently, we show how this verification task can be approached by three-valued spotlight abstraction.

Spotlight abstraction involves the partition of the processes of the system into a *spotlight* and a *shade*. Predicate abstraction is applied to the spotlight, while the shade processes are abstracted away by summarising them in one approximative component. The state space of the resulting abstract system can be straightforwardly modelled as a (pure) three-valued Kripke structure. In our current verification task, the relevant process for the property of interest is $Proc_1$,

which we put into the spotlight: $Spot(Proc) = \{Proc_1\}$, whereas the remaining system is for now kept in the shade: $Shade(Proc) = \{Proc_2, \ldots, Proc_n\}$. Next, a set of so-called *spotlight predicates* over the system variables is selected, here we choose $Spot(Pred) = \{progress, (v_1 > 0)\}$. By applying three-valued predicate abstraction to the spotlight processes, we obtain an abstract process $Proc_1^a$ with the same control flow as $Proc_1$ but with operations abstracted over $Spot(Pred)$. The processes in the shade are summarised to one approximative process $Proc_{Shade}$. Due to the loss of information about the shade, $Proc_{Shade}$ might set predicates over shared variables to the value $\perp$. Our abstract system now looks as follows: $Sys^a = Proc_1^a \parallel Proc_{Shade}$. The state space of Sys^a can be modelled as a pure three-valued Kripke structure over $AP = Spot(Pred) \cup \{(pc_i = j) \mid Proc_i \in Spot(Proc), j \in Loc_i\}$ where $(pc_i = j)$ refers to the program counter of $Proc_i$, and each definite model checking result obtained for this structure can be transferred to the concrete system [12]. A three-valued Kripke structure K corresponding to Sys^a is depicted in part (a) of the figure below. For simplicity, we only show the program counter predicates that are currently *true*.

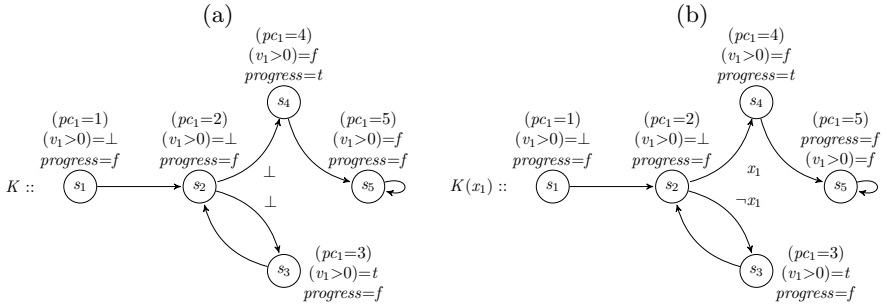

Note that the control flow of spotlight processes is always preserved under spotlight abstraction. Hence, each transition of K associated with the spotlight matches with a specific operation of the spotlight process $Proc_1$. For K and its set of atomic predicates $AP = \{progress, (v_1 > 0)\} \cup \{(pc_1 = j) \mid j \in Loc_1\}$ we can formalise our property of interest as the LTL formula $\mathbf{GF}progress$ and then apply standard three-valued model checking, i.e. check $[K, s_1 \models \mathbf{GF}progress]$. The current abstraction is not precise enough for a definite result in verification. Since there exist processes in the shade that operate on the shared variable v_1, the value of the predicate $(v_1 > 0)$ in the states s_1 and s_2 is $\perp$. Thus, it is also unknown whether the body of the *while*-loop can be executed via the transition (s_2, s_3), or whether the loop can be eventually left via (s_2, s_4). The automatic abstraction refinement procedure introduced in [17] would now iteratively shift processes from the shade to the spotlight until it can be definitively shown *which* branch of the *while*-loop can be actually taken. However, due to transitive dependencies – $Proc_2$ modifies v_1, but in turn depends on $v_2, \ldots, v_k$ which may be modified by other shade processes as well – such a refinement can be exceedingly costly or can even lead to a failure of verification because of state explosion. A closer look at our simple example structure tell us that, regardless of which branch of the loop will be ever taken, *progress* will never hold repeatedly. Hence, the evaluation of $\mathbf{GF}progress$ on K should yield *false*. However, the standard

three-valued LTL semantics (compare Section 2) does not allow us to draw this conclusion. In the following we will see that automated *parameterisation* can give us the necessary precision for a definite verification result – at considerably less cost than classical abstraction refinement.

4.2 Parameterisation of Three-Valued Abstractions

As we just have seen, $[K, s_1 \models \mathbf{GF}\mathit{progress}]$ yields $\bot$. Nevertheless, a $\bot$-result in 3MC always comes along with an *unconfirmed counterexample* – a potential error path in the Kripke structure with some *unknown* transitions or predicates. Four our running example the path $\pi = s_1 s_2 s_4 s_5 s_5 \ldots$ is an unconfirmed counterexample. Such a path is typically used for *counterexample-guided abstraction refinement* (CEGAR) [4]: In our case, the $\bot$-transition (s_2, s_4) would be identified as the reason for uncertainty, and shade processes that modify the *if*-condition $(v_1 > 0)$ associated with (s_2, s_4) would be iteratively shifted to the spotlight. Now we will show that counterexamples can also be exploited for the parameterisation of three-valued Kripke structures. We first illustrate parameterisation based on our running example and then provide the general rules for it.

Our method detects that the reason for uncertainty, the $\bot$-transition (s_2, s_4) along π, is associated with a *complementary branch* in the original system: a branch of the control flow of a single process with complementary branching conditions – here $(v_1 > 0)$ and $\neg(v_1 > 0)$. Instead of applying classical CEGAR, a fresh boolean parameter x_1 is introduced and the transition is parameterised as follows: $R(s_2, s_4) := x_1$. Next, the complementary transition (s_2, s_3) is identified and parameterised by $R(s_2, s_3) := \neg x_1$. The corresponding parameterised three-valued Kripke structure $K(x_1)$ is depicted in part (b) of the figure on the previous page. Applying parameterised three-valued model checking, i.e. verifying $[K(x_1), s_1 \models \mathbf{GF}\mathit{progress}]$ immediately returns *false*. Thus, for our running example a definite result in verification only requires the introduction of a single parameter and the consideration of the two instantiations $K(\mathit{true})$ and $K(\mathit{false})$ of $K(x_1)$. In contrast, a corresponding pure three-valued approach would require a large number of additional refinement steps and thus would most likely fail due to state explosion. Also the application of the computationally more expensive GMC would not be successful, since it cannot establish the complementary relation between (s_2, s_4) and (s_2, s_3). The following rule generalises the parameterisation of complementary branches in three-valued Kripke structures.

Rule I (Parameterisation of Complementary Branch Transitions). *Let $Sys = \|_{i=1}^{n} Proc_i$ be a concurrent system and $Spot = Spot(Proc) \cup Spot(Pred)$ be a spotlight abstraction for Sys. Let K be a three-valued KS over $AP = Spot(Pred) \cup \{(pc_i = j) \mid Proc_i \in Spot(Proc) \wedge j \in Loc_i\}$ that models the abstract state space corresponding to Sys and $Spot$, and let s_1 be a state of K. Moreover, let ψ be a safety or liveness LTL formula and checking $[K, s_1 \models \psi]$ yields $\bot$. Let π be the unconfirmed counterexample returned by model checking which runs through a finite number of different transitions. The transitions of K can be parameterised as follows:*

For each transition (s, s') along π with $R(s, s') = \bot$, check if (s, s') is part of a complementary branch, i.e.: (s, s') is associated with a guard operation assume(e) of a spotlight process $Proc_i$, where e is a boolean expression – and moreover, there exists a state s'' such that (s, s'') is associated with a complementary guard operation assume$(\neg e)$ of $Proc_i$. Then introduce a fresh parameter x_j and set $R(s, s') = x_j$ and $R(s, s'') = \neg x_j$.

This rule allows to parameterise complementary branches (e.g. *if*- or *while*-operations) in three-valued abstractions. As we have seen in our running example, this can lead to substantial savings in the number of necessary refinement steps for a definite result in verification. In fact, any verification task where the property of interest turns out to be independent from certain branches can profit from such a parameterisation in a similar manner. At the end of this section we will present a theorem which states that the application of Rule I leads to sound abstractions of concurrent systems. Beforehand, we introduce another rule that allows the parameterisation of *predicates* in three-valued abstractions.

In order to illustrate how the parameterisation of predicates works, we consider a second example, the concurrent system *Sys* depicted below. Our property of interest is now *mutual exclusion*, i.e. whether the flag variables $flag_1$ and $flag_2$ are never *true* at the same time.

$$v_1, \ldots, v_k : \textbf{integer};$$

$$flag_1, flag_2, init : \textbf{boolean where } flag_1 = false, flag_2 = false, init = false;$$

$$Proc_1 :: \begin{bmatrix} 1: & flag_1 := f(v_1, \ldots, v_k) \\ 2: & init := true \\ 3: & flag_1 := \neg flag_2 \\ 4: & [\ \ldots \quad] \end{bmatrix} \ \| \ Proc_2 :: \begin{bmatrix} 1: & flag_2 := false \\ 2: & \textbf{await}(init) \\ 3: & flag_2 := \neg flag_1 \\ 4: & [\ \ldots \quad] \end{bmatrix} \ \| \ \ldots \ \| \ Proc_n$$

Applying three-valued spotlight abstraction with classical refinement yields the following spotlight after a number of iterations: $Spot(Proc) = \{Proc_1, Proc_2\}$ and $Spot(Pred) = \{flag_1, flag_2, init\}$. Next, a corresponding pure three-valued Kripke structure K over $AP = \{flag_1, flag_2, init\} \cup \{(pc_i = j) \mid Proc_i \in Spot(Proc) \wedge j \in Loc_i\}$ is constructed, and the mutual exclusion property formalised by the safety LTL formula $\mathbf{G}\neg(flag_1 \wedge flag_2)$ is checked for K. Model checking returns *unknown*, since the assignment to $flag_1$ at location 1 of $Proc_1$ depends on the shared variables $v_1, \ldots, v_k$ which are potentially modified by a large number of processes that are currently in the shade. Thus, with classical abstraction refinement we have to expect a large number of further refinement steps necessary for a definite result in verification: Predicates over the variables $v_1, \ldots, v_k$ as well as processes modifying these variables have to be drawn into the spotlight. Nevertheless, the model checking run based on the current spotlight also returns the unconfirmed counterexample π depicted in part (a) of the figure below.

(a) (b)

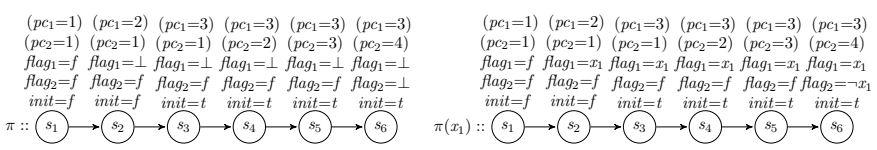

The reason for uncertainty is the reachable state s_6 where flag_1 and flag_2 are both $\bot$. The predicate flag_1 is set to $\bot$ by transition (s_1, s_2), since there are not enough predicates and processes in the spotlight in order to abstract the associated operation $\mathit{flag}_1 := f(v_1, \ldots, v_k)$ properly. The predicate flag_2 is set to $\bot$ by (s_5, s_6) because the associated operation $\mathit{flag}_2 := \neg \mathit{flag}_1$ modifies this predicate in relation to the already *unknown* predicate flag_1. In our simple example it is easy to see that flag_1 and flag_2 must have *complementary* values in state s_6 – which would rule out the unconfirmed counterexample π. However, this fact cannot be captured by pure three-valued abstraction since it does not allow to establish connections between predicates that are associated with the value $\bot$.

Our concept of parameterisation enables us to establish such connections. For our running example we proceed as follows: We backtrack to the state s_2 where flag_1 was initially associated with $\bot$. Next, we introduce a fresh parameter x_1 and set $L(s_2, \mathit{flag}_1) := x_1$. Based on the operations associated with the succeeding transitions along π we update the labellings of the states s_3 to s_6. As a consequence, we now can capture that flag_1 constantly keeps the value x_1 along π, flag_2 keeps the value *false* until s_5, and in particular, flag_1 and flag_2 have complementary values in s_6. The resulting path $\pi(x_1)$, which is depicted in part (b) on the previous page, is no longer an unconfirmed counterexample. Thus, checking $\mathbf{G}\neg(\mathit{flag}_1 \wedge \mathit{flag}_2)$ on a corresponding parameterised Kripke structure $K(x_1)$ will immediately return that no counterexample exists, i.e. that the property is satisfied for the modelled system. Again we have seen that parameterisation – here with regard to predicates – can lead to substantial savings in the number of necessary refinement steps for a definite result in verification. The following rule generalises the parameterisation of predicates in three-valued abstractions.

Rule II (Parameterisation of Predicates along Counterexamples). *Let Sys, Spot, K, s_1 and AP be as in Rule I. Moreover, let $\psi = \mathbf{G}\neg(\bigwedge_{i=1}^{m} p_i)$ be a safety LTL formula with $\{p_1, \ldots, p_m\} \subseteq \mathit{Spot}(\mathit{Pred})$ and model checking $[K, s_1 \models \psi]$ yields $\bot$. Let $\pi = s_1 \ldots s_k$ be the unconfirmed counterexample returned by model checking which is a path prefix that ends in a state s_k where all predicates from $\{p_1, \ldots, p_m\}$ are associated with either the value $\bot$ or true. K can be parameterised along π according to the following procedure:*

for $s := s_1$ *to* s_k **do**
  for *each* $p_i \in \{p_1, \ldots, p_m\}$ *with* $L(s_k, p_i) = \bot$ **do**
    if $L(s, p_i) = \bot$ **then**
      if $s = s_1$, *i.e. s is the initial state* **then**
        introduce a fresh parameter x_j and set $L(s, p_i) := x_j$
      else
        let s' be the direct predecessor of s along π, and let *op* be the operation associated with the transition (s', s)
        if *op is not associated with a process in $\mathit{Spot}(\mathit{Proc})$ or none of the atomic predicates occurring in the weakest precondition*[1] $wp_{op}(p_i)$ *are contained in $\mathit{Spot}(\mathit{Pred})$* **then**
          introduce a fresh parameter x_j and set $L(s, p_i) := x_j$
        else
          set $L(s, p_i) :=$
          $wp_{op}(p_i)\,[p/L(s', p) \mid p \in \mathit{Spot}(\mathit{Pred})]\,[p/\bot \mid p \notin \mathit{Spot}(\mathit{Pred})]$,
          i.e. update $L(s, p_i)$ wrt. parameterisations in predecessor s'

[1] Let $op = \mathit{assume}(e) : x_1 := e_1, \ldots, x_m := e_m$ then $wp_{op}(p) = e \wedge p[x_1/e_1, \ldots, x_m/e_m]$.

Parameterisation of predicates is applied in a similar way for model checking liveness formulae, i.e. $[K, s_1 \models \mathbf{GF}(\bigvee_{i=1}^{m} p_i)]$ *with* $\{p_1, \ldots, p_m\} \subseteq Spot(Pred)$. *In case of an unknown result, the model checker additionally returns an unconfirmed counterexample* π *of the form* $(s_1 \ldots s_{l-1}) \circ (s_l \ldots s_k)^\omega$ *and in all states* $s_l \ldots s_k$ *each predicate from* $\{p_1, \ldots, p_m\}$ *is associated with either the value* $\perp$ *or false. The finite prefix* $(s_1 \ldots s_{l-1})$ *of* π *is then parameterised in the same manner as in the case of model checking safety formulae.*

The following theorem establishes the soundness, with respect to the information ordering $\leq_{\mathbb{K}_3}$ (compare Section 2), of parameterised three-valued model checking, provided that parameterisation is applied according to Rule I and II.

Theorem 1. *Let Sys and Spot be as before. Let* K *over* AP *be a two-valued KS modelling the concrete state space of Sys and let* $K^\perp$ *over* $AP^\perp = Spot(Pred) \cup \{(pc_i = j) \mid Proc_i \in Spot(Proc) \wedge j \in Loc_i\}$ *with* $AP^\perp \subseteq AP$ *be a pure three-valued KS modelling the abstract state space corresponding to Spot. Moreover, let* s_1 *and* $s_1^\perp$ *be states representing the initial configuration of Sys in* K *resp.* $K^\perp$. *Then for any parameterisation* $K^\perp(\overset{m}{x})$ *of* $K^\perp$ *obtained by applying the rules I and II, and for any safety or liveness LTL formula* ψ[2] *over* $AP^\perp$ *the following holds:*

$$[K^\perp(\overset{m}{x}), s_1^\perp \models \psi] \quad \leq_{\mathbb{K}_3} \quad [K, s_1 \models \psi]$$

Proof. See `http://www.cs.up.ac.za/cs/ntimm/proof.pdf`

Hence, every definite result in verification obtained for $[K^\perp(\overset{m}{x}), s_1^\perp \models \psi]$ can be directly transferred to the concrete system modelled by K, whereas an *unknown* result for $[K^\perp(\overset{m}{x}), s_1^\perp \models \psi]$ tells us that further abstraction refinement or parameterisation of $K^\perp(\overset{m}{x})$ is required. In the next section, we will show how we have implemented the application of the parameterisation rules within an automatic abstraction refinement procedure for the verification of concurrent systems and how verification can benefit from our parameterisation approach.

5 Automatic Counterexample-Guided Refinement and Parameterisation

We have prototypically implemented a verification framework for concurrent systems based on spotlight abstraction with counterexample-guided refinement and parameterisation. Our framework 3Spot works on top of the three-valued symbolic model checker χChek [5]. 3Spot takes a concurrent system *Sys* over a variable set *Var* and a safety or liveness temporal logic formula ψ over *Sys* as input. The initial spotlight *Spot* is defined by the processes that are referenced in ψ and the atomic predicates over *Var* that are subformulae of ψ. Next, a parameterised three-valued Kripke structure $K^\perp(\overset{m}{x}) = (S, R, L, \mathbb{F})$ corresponding to *Sys* and *Spot* is constructed with a state $s_1 \in S$ representing the initial configuration of *Sys*. The parameter tuple $(\overset{m}{x})$ of $K^\perp(\overset{m}{x})$ is initially empty. In order to check $[K^\perp(\overset{m}{x}), s_1 \models \psi]$, the following procedure is executed:

[2] ψ is either of the form $\mathbf{G}\neg(\bigwedge_{i=1}^{m} p_i)$ or $\mathbf{GF}(\bigvee_{i=1}^{m} p_i)$ with $\{p_1, \ldots, p_m\} \subseteq AP^\perp$.

1. **check** $[K^{\perp}(\overset{m}{a}), s_1 \models \psi]$ for all valuations $(\overset{m}{a}) \in \{t,f\}^m$
   if $\forall(\overset{m}{a}) \in \{t,f\}^m : [K^{\perp}(\overset{m}{a}), s_1 \models \psi] = t$ **or** $\forall(\overset{m}{a}) \in \{t,f\}^m : [K^{\perp}(\overset{m}{a}), s_1 \models \psi] = f$ **then**
   property ψ is successfully proved resp. disproved for the concurrent system Sys; stop
   if $\forall(\overset{m}{a}) \in \{t,f\}^m : [K^{\perp}(\overset{m}{a}), s_1 \models \psi] \in \{\perp, t\}$ **or** $\forall(\overset{m}{a}) \in \{t,f\}^m : [K^{\perp}(\overset{m}{a}), s_1 \models \psi] \in \{\perp, f\}$ **then**
   still some *unknown* results; further refinement or parameterisation required; go to 2.
   if $\exists(\overset{m}{a}) \in \{t,f\}^m : [K^{\perp}(\overset{m}{a}), s_1 \models \psi] = t$ **and** $\exists(\overset{m}{a}) \in \{t,f\}^m : [K^{\perp}(\overset{m}{a}), s_1 \models \psi] = f$ **then**
   current parameterisation not expedient; revoke last parameterisation; go to 2.
2. **for** each valuation $(\overset{m}{a}) \in \{t,f\}^m$ with $[K^{\perp}(\overset{m}{a}), s_1 \models \psi] = \perp$ **do**
   generate unconfirmed counterexample $\pi^{\perp}$ for $[K^{\perp}(\overset{m}{a}), s_1 \models \psi]$
   select unconfirmed counterexample $\pi^{\perp}$ with the fewest *unknown* transitions and predicates
   if Rule I is applicable along $\pi^{\perp}$ **then**
   apply Rule I to the corresponding branch in $K^{\perp}(\overset{m}{x})$
   else if Rule II is applicable along $\pi^{\perp}$ **then**
   apply Rule II to the corresponding path prefix in $K^{\perp}(\overset{m}{x})$
   else
   determine cause of indefinite result along $\pi^{\perp}$ and derive corresponding refinement candidate r (see our previous work [17] for an example technique for deriving refinement candidates from unconfirmed counterexamples), which can be a shade process or a predicate; add r to *Spot*
   if r is a predicate **then**
   revoke parameterisation for parameterised branches in $K^{\perp}(\overset{m}{x})$ where the value of r affects the branching condition
   update $K^{\perp}(\overset{m}{x})$ according to changes in 2. and go to 1.

Hence, the procedure terminates if for all instantiations of the current parameterised Kripke structure the same definite result in verification can be obtained. If model checking yields *true* for some instantiations and *false* for others, the last parameterisation step was not expedient: The property of interest is then obviously not independent from the most recent parameterisation. Thus, this step is revoked, which also includes that the same parameterisation will not be admissible in future iterations. In case model checking returns *unknown* for some instantiations, the abstraction has to be further parameterised or refined based on unconfirmed counterexamples obtained for these instantiations. For this purpose we always apply Rule I or II if possible, or use classical refinement (see our previous work [17]) otherwise. Adding a new predicate p to the abstraction may affect parameterised branches: An abstract state s that is the starting point of a complementary branch may be split into two new states s_a and s_b with $L(s_a, p) = \mathit{true}$ and $L(s_b, p) = \mathit{false}$. Thus, in the general case, the parameterisation of the complementary branch starting in s has to be revoked. However, if the branch condition is independent from the value of p then the parameterisation can be kept. Alternatives to the revocation of parameterisations are: Keeping the parameterisation for only one state, either s_a or s_b. Or, introducing a fresh parameter x_j for the second branch starting in s_b. Each iteration ends with the update of the parameterised three-valued Kripke structure according to new parameterisations or additional refinements. In case a new predicate has been added to the abstraction, this update also involves the recalculation of the parameterisation of predicates (compare last step of Rule II).

So far, parameterisation resp. refinement is performed based on the unconfirmed counterexample with the fewest *unknown* transitions and predicates. The intention behind this is to minimise the expected effort to confirm or eliminate the counterexample. Moreover, the attempt to apply the parameterisation rules

or classical refinement is so far always conducted in the fixed order *Rule I, Rule II, refinement*. In the future, we intend to use heuristic guidance for selecting the unconfirmed counterexample and for deciding which rule application or which refinement step is currently most promising in order to achieve a definite result in verification within a small number of iterations. Similar to our previous work on heuristics for pure refinement [17], we plan to base this heuristic approach on the structure of the underlying concurrent system, i.e. on the variable dependencies between the processes of the system.

In preliminary experiments, we applied our procedure to multiple-resource allocation systems[3] with up to 25 processes and 140 variable dependencies, and we checked safety as well as liveness properties. We compared verification under the pure three-valued approach (which has proven to be generally successful for concurrent systems in [17, 14, 15]) with verification under our novel approach with parameterisation. In several cases where the pure three-valued approach failed due to an out-of-memory exception, our new technique was capable of returning a definite verification result. The additional computations for parameterisation particularly paid off when the property of interest turned out to be independent from certain branches in the system, and the costs for concretising these branches via classical refinement were high. In fact, such cases are very common for systems with many *if*-, *while*-, and similar operations. We also observed verification tasks (primarily where the system only exhibited very few branches, or where the property was dependent on most of the branches) that did not profit from the application of parameterisation rules. Here verification under the new approach was slower but did not fail, since parameterisation only increases the number of checks per iteration, but not the size of the abstraction (spotlight processes and predicates). Thus, so far it is a good strategy to apply the pure three-valued approach first and in case of failure the approach with parameterisation subsequently. Nevertheless, with our intended heuristic approach, we aim at directly discovering the best possible combination of refinement and parameterisation for each verification task. A more extensive experimental evaluation of such an enhanced approach is also planned as future work.

6 Related Work

Or research is situated in the field of model checking temporal logic properties on partial system models. The idea of evaluating temporal logic formulae on three-valued Kripke structures was initially proposed in [3] and is now established under the name *three-valued model checking* (3MC). Our new concept *parameterised three-valued model checking* (PMC) is an extension of 3MC. In our approach, unknown parts of the modelled system cannot only be represented by the constant $\bot$, but also by expressions over boolean parameters. The evaluation of temporal logic formulae is then performed for each possible parameter instantiation. The idea of considering possible instantiations resp. concretisations of a partial model is adopted from *generalised model checking* (GMC) [7]. In contrast to the concretisations in GMC, our instantiations only affect parameters

[3] A detailed description of these systems can be found in [14].

but do not concern the constant $\perp$. Moreover, our instantiations are always of the same size as the partial model, whereas the concretisations in GMC can be exponentially larger. Neither 3MC nor GMC offer a concept for drawing connections between unknown parts. While 3MC and GMC are general concepts for the verification of partial models, our approach is application-oriented and takes advantage from the consideration of the system structure when applying the parameterisation rules within our automated verification procedure.

Another work related to ours is that of Herbstritt et al. [9] who combine three-valued logic and quantified boolean parameters for representing unspecified parts of a hardware model with different precision. Their technique is geared towards equivalence checking of circuits. In contrast to our approach, [9] do not introduce a concept for establishing connections between parameters in the model. Moreover, the decision for modelling an unspecified part via the third truth value $\perp$ or via a boolean parameter has to be done by hand and not based on automatable rules. [9] encode their hardware verification tasks as bounded model checking problems that can be efficiently solved via SAT/QBF-solvers. The definition of such encodings for our parameterised three-valued model checking is another interesting direction for future research. A similar approach to the verification of hardware circuits, but in the context of BDD-based symbolic model checking was introduced in [11]. Their method supports the verification of full CTL properties based on models with a flexible representation of unknowns. This approach necessitates the manual selection of the type of modelling unknown parts. Establishing logical relations between parameters is not possible here.

7 Conclusion

We developed a concept for modelling unknown parts of an abstract software system with different types of approximation: In our parameterised three-valued Kripke structures the loss of information about a predicate or a transition can be either represented by the constant $\perp$ or by an expression over boolean parameters. A novel feature of our modelling approach is that it allows for establishing logical connections between *unknown* parameters, like equality or complementarity – and thus, to preserve more details under abstraction that can be crucial for the success and efficiency of verification. We introduced temporal logic model checking for parameterised three-valued Kripke structures and showed that this method is sound if the models are constructed with regard to parameterisation rules that we defined. These rules take the branching structure and the program code of the modelled system into account and arrange the connections between parameters in the model. We then presented an automatic verification procedure based on iterative abstraction refinement and parameterisation. For several verification tasks, particularly for verifying systems with many conditional branches, our new approach with parameterisation can significantly outperform verification based on classical modelling techniques that are not capable of characterising connections between unknown parts. We are convinced that our concept for parameterisation can be easily and effectively adapted to other types of systems and verification tasks, which we intend to investigate in our future research.

References

1. Alfaro, L., Roy, P.: Solving games via three-valued abstraction refinement. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007, LNCS, vol. 4703, pp. 74–89. Springer-Verlag Berlin Heidelberg (2007)
2. Ball, T., Majumdar, R., Millstein, T., Rajamani, S.K.: Automatic predicate abstraction of C programs. In: ACM SIGPLAN 2001. pp. 203–213. PLDI '01, ACM, New York, NY, USA (2001)
3. Bruns, G., Godefroid, P.: Model checking partial state spaces with 3-valued temporal logics. In: Halbwachs, N., Peled, D. (eds.) CAV 1999. pp. 274–287. LNCS, Springer-Verlag Berlin Heidelberg, London, UK (1999)
4. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000, LNCS, vol. 1855, pp. 154–169. Springer-Verlag Berlin Heidelberg (2000)
5. Easterbrook, S.M., Chechik, M., Devereux, B., Gurfinkel, A., Lai, A.Y.C., Petrovykh, V., Tafliovich, A., Thompson-Walsh, C.: χChek: A model checker for multi-valued reasoning. In: ICSE 2003. pp. 804–805 (2003)
6. Fitting, M.: Kleene's three valued logics and their children. Fundamenta Informaticae 20(1-3), 113–131 (Mar 1994)
7. Godefroid, P., Piterman, N.: LTL generalized model checking revisited. In: Jones, N.D., Mueller-Olm, M. (eds.) VMCAI 2009, LNCS, vol. 5403, pp. 89–104. Springer Berlin Heidelberg (2009)
8. Grumberg, O.: 2-valued and 3-valued abstraction-refinement in model checking. In: Logics and Languages for Reliability and Security, pp. 105–128. IOS Press, Incorporated (2010)
9. Herbstritt, M., Becker, B.: On combining 01X-logic and QBF. In: Moreno Diaz, R., Pichler, F., Quesada Arencibia, A. (eds.) Comp. Aided Systems Theory - EUROCAST 2007, LNCS, vol. 4739, pp. 531–538. Springer Berlin Heidelberg (2007)
10. Katoen, J.P., Klink, D., Leucker, M., Wolf, V.: Three-valued abstraction for probabilistic systems. Logic and Algebraic Programming 81(4), 356 – 389 (2012)
11. Nopper, T., Scholl, C.: Symbolic model checking for incomplete designs with flexible modeling of unknowns. IEEE Trans. Computers 62(6), 1234–1254 (2013)
12. Schrieb, J., Wehrheim, H., Wonisch, D.: Three-valued spotlight abstractions. In: Cavalcanti, A., Dams, D.R. (eds.) FM 2009: Formal Methods, LNCS, vol. 5850, pp. 106–122. Springer-Verlag Berlin Heidelberg (2009)
13. Shoham, S., Grumberg, O.: 3-valued abstraction: More precision at less cost. Information and Computation 206(11), 1313 – 1333 (2008)
14. Timm, N.: Three-Valued Abstraction and Heuristic-Guided Refinement for Verifying Concurrent Systems. Phd thesis, University of Paderborn (2013)
15. Timm, N.: Spotlight abstraction with shade clustering – automatic verification of parameterised systems. In: 8th International Symposium on Theoretical Aspects of Software Engineering, IEEE Computer Society (to appear) (2014)
16. Timm, N., Wehrheim, H.: On symmetries and spotlights – verifying parameterised systems. In: Dong, J., Zhu, H. (eds.) ICFEM 2010, LNCS, vol. 6447, pp. 534–548. Springer, Heidelberg (2010)
17. Timm, N., Wehrheim, H., Czech, M.: Heuristic-guided abstraction refinement for concurrent systems. In: Aoki, T., Taguchi, K. (eds.) ICFEM 2012, LNCS, vol. 7635, pp. 348–363. Springer Berlin Heidelberg (2012)
18. Wei, O., Gurfinkel, A., Chechik, M.: On the consistency, expressiveness, and precision of partial modeling formalisms. Information and Computation 209(1), 20 – 47 (2011)

A Probabilistic Model Checking Analysis of a Realistic Vehicular Networks Mobility Model

Bruno Ferreira, Fernando A. F. Braz, and Sérgio V. A. Campos

Department of Computer Science, Federal University of Minas Gerais
Av. Antônio Carlos, 6627, Pampulha, 30123-970 Belo Horizonte, Brazil
{bruno.ferreira,fbraz,scampos}@dcc.ufmg.br

Abstract. Vehicular Ad-Hoc Networks (VANET) are a special type of network where its nodes are vehicles that move according to specific patterns. This network is based on wireless communication, presenting new challenges, such as how it will be tested in realistic scenarios. Currently, simulations are widely used. However, they have limitations, such as local minima. Another approach is model checking, which has been used in only a few studies, often overlooking mobility and signal propagation issues. This work provides a realistic mobility model using probabilistic model checking to describe an overtake scenario involving three vehicles in a short distance. Our analysis has shown 98% of accident chance in this situation. However, the main result is providing an example to represent the mobility aspect which can be connected with other models such as signal propagation and the network itself. Therefore, VANETs can now be tested using methods closer to the reality.

Keywords: Model Checking, Vehicular Ad-Hoc Networks, Mobility

1 Introduction

Intelligent Traffic Systems (ITS) are a response to reduce the number of traffic accidents, the cost of transportation and the volume of CO_2 emissions [11]. These systems make intensive use of communication among vehicles, which is possible using Vehicular Ad-Hoc Networks (VANETs), a particular class of Mobile Ad-Hoc Networks (MANETs). VANETs are distributed and self-organized communication networks, characterized by their high speed and mobility, which brings several challenges to the academic community [13].

Current research in this field frequently analyzes the behavior of VANETs using simulators. However, simulation methods examine only a subset of possible scenarios, which can lead to an incomplete – or even worse, an incorrect – analysis [14]. Furthermore, works such as [2] and [4] have reported that VANET simulators, despite their constant evolution, have not reached an ideal point, because they need to integrate the mobility of the nodes, the communication protocols (network model) and the signals propagation.

A complementary approach to simulation is the use of probabilistic model checking (PMC) [9, 19]. PMC is a technique for the automatic analysis of systems, which verifies properties in probabilistic logic by exhaustively enumerating

all reachable states. PMC can answer questions such as "What is the probability of the occurrence of a certain event?". This approach is ideal for dynamic and stochastic systems, such as VANETs. PMC verification is performed by (1) specifying what are the properties that the system must obey, (2) constructing the formal model of the system, which should capture all the essential properties and (3) finally, running the verifier to validate the specified properties.

Verification techniques can be useful to assess the efficiency and correctness of MANETs. The results obtained can be used to improve a wide range of systems. Despite its benefits, model checking is rarely used in VANETs. Also, the few studies (e.g. [5] and [21]) do not address uncertainty caused by the dynamism of the nodes. Thus, the non-determinism of the message delivery caused by the mobility of vehicles is not being represented, which is an underlying factor in VANETs. [20] uses simulation of Markov chains to represent planned trajectories of autonomous vehicles. The tool which we have used for analysis also represents its model with this technique, however, it uses a formal approach, finding exact probabilities and estimates, besides, other resources such as multi-terminal binary decision diagrams are used [19] and our work benefits from these features.

It is important to verify networks considering not only the network itself, but also its additional functionalists. Thus, it is often necessary to model the communication and other important system components [8]. Therefore, building complete models considering the traffic flow, network and radio propagation are necessary and rarely explored in model checking. We have proposed the first step for completely modeling VANETs presenting a motion aspect which will be coupled with the traditional network analysis.

Nevertheless, this work has the objective of representing mobility models in VANETs using PMC. The proposed model follows practices and concepts already used in simulation methods to model an overtake situation involving three vehicles. However, it uses the benefits of automatic and exhaustive verification provided by PMC. Thus, the application of model checking in VANETs can be extended in the future to describe network and mobility models.

We have used PRISM, a probabilistic model checker for formal modelling and analysis. This tool can represent systems that exhibit random or probabilistic behaviour. It has been used to analyze many different application domains from communication protocols to biological systems [18]. We have modeled an overtake scenario involving three vehicles in a short distance. The model shows that there is a huge chance of an accident (98% in some scenarios), however counter-examples to a safe overtake are presented.

This paper is organized as follows: Section 2 presents important concepts of VANET analysis; PMC is defined in Section 3; Section 4 shows our mobility model; Section 5 discusses the results of the model; finally, conclusions and future works are presented in Section 6.

2 VANET Analysis

In order to validate the effectiveness of Intelligent Traffic Systems, it is necessary to evaluate their performance and communication protocols in real test environments. However, there are logistic difficulties, economic questions and technological limitations which make simulations a good choice for testing and validation of these protocols. The fields of computer networks and traffic engineering make extensive use of simulators. There are long established software such as NS-2 (The Network Simulator)[1] and SUMO (Simulation of Urban Mobility)[2]. Since the introduction of vehicular networks, the integration of these two fields has recently become necessary [14].

This integration is required due to inherent features of the strong coupling between communication and mobility in VANETs. Communication modifies mobility patterns, on the other hand, correct message reception is affected by vehicular movement. However, three distinct aspects must work together in order to achieve realistic tests [4]: (1) **Mobility Models** represent the vehicle movement, including mobility patterns and the interaction between vehicles (e.g. crossroad control); (2) **Network Models** describe the data exchanged between vehicles, including MAC, routing and superior protocol layers; (3) **Signal Propagation Models** reproduce the environment modeling involving fixed and mobile obstacles during the communication. For further details on these mobility and signal propagation techniques, we refer to [13] and [17], respectively.

Mobility models, the main subject of this work, can be described in two points [12]: (1) **Freedom of movement**, responsible for describing the motion constraint to each vehicle. These representations have been improved from simplified models such Manhattan grid [3] to real world maps (e.g. [7] and [22]) and (2) **Interaction among vehicles** which modeling the behavior of a vehicle that is a direct consequence of the interaction with the other vehicles on the road. This includes microscopic aspects, such as lane changing and decreasing/increasing the speed due to the surrounding traffic.

Regarding this microscopic implementation, Car Following Models (CFMs) are the most used type of driver model. CFMs usually represent time, position, speed, and acceleration as continuous functions. However, CFMs have been extended to include discrete formulations [14]. Commonly used models are (as described by [15]): the cellular automata models, follow-the-leader models and intelligent driver model (IDM). The next subsections describe two CFM models chosen for their simplicity, efficiency and realism.

2.1 Intelligent Driver Model

The Intelligent Driver Model (IDM) shows a crash-free collective dynamic, exhibits controllable stability properties, and implements a braking strategy with smooth transitions between acceleration and deceleration behavior [16]. The IDM

[1] NS-2 . `http://www.isi.edu/nsnam/ns/`. Access date: September 10, 2014
[2] SUMO. `http://sumo.sourceforge.net/`. Access date: September 10, 2014

acceleration is a continuous function incorporating different driving modes for all velocities of freeway and city traffic. The distance s (bumper-to-bumper) to the leading vehicle is given by $s = x_l - x - e$, where x_l and x are the coordinates and e is the extent of vehicle. IDM also takes into account the velocity difference (approaching rate) to the leading vehicle, given by $\Delta v = v - v_l$. The IDM acceleration function is given by the Equations 1 and 2.

$$a_{IDM}(s, v, \Delta v) = a \left[1 - \left(\frac{v}{v_0} \right)^{\delta} - \left(\frac{s^*(v, \Delta v)}{s} \right)^2 \right] \tag{1}$$

$$s^*(v, \Delta v) = s_0 + vT + \frac{v \Delta v}{2\sqrt{ab}} \tag{2}$$

This expression combines the free-road acceleration strategy, given by:

$$a_{free}(v) = a[1 - (v/v_0)^{\delta}]$$

with a deceleration strategy, given by:

$$a_{brake}(s, v, \Delta v) = -a(s^*/s)^2$$

The deceleration strategy becomes relevant when the gap to the leading vehicle is not significantly larger than the "desired (safe) gap", given by $s^*(v, \Delta v)$. The free acceleration is denoted by the desired speed v_0, the maximum acceleration is a, and the exponent δ indicates how the acceleration decreases with velocity ($\delta = 1$ corresponds to a linear decrease, while $\delta \to \infty$ denotes a constant acceleration).

The effective minimum gap s^* is composed of the minimum distance s_0 (which is relevant for low velocities only), the velocity dependent distance vT, which corresponds to following the leading vehicle with a constant desired time gap T, and a dynamic contribution which is only active in non-stationary traffic corresponding to situations in which $\Delta v \neq 0$. This latter contribution implements an "intelligent" driving behavior that, in normal situations, limits braking decelerations to a comfortable deceleration b. In critical situations, however, the IDM deceleration becomes significantly higher, making the IDM collision-free [24]. The IDM parameters v_0, T, s_0, a and b are shown in Table 1.

Table 1. Parameters of the Intelligent Driver Model. Adapted– [16]

Parameter	Car	Truck
Desired speed v_0	120 km/h	85 km/h
Free acceleration exponent δ	4	4
Desired time gap T	1.5	2.0
Jam distance s_0	2.0	4.0
Maximum acceleration a	1.4 m/s^2	0.7 m/s^2
Desired deceleration b	2.0 m/s^2	2.0 m/s^2
Changing threshold Δ_{th}	0.1 m/s^2	0.1 m/s^2

Calculating the acceleration at a time t, the new position and speed or deceleration distance can be given by traditional kinematics' equations.

2.2 Minimizing Overall Braking Induced by Lane Change

A general model to represent lane-changing rules was proposed by [10]. The model is called Minimizing Overall Braking Induced by Lane Change (MOBIL). The utility and risk associated of a given lane are determined in terms of longitudinal accelerations calculated by microscopic car-following models as IDM. The previous vehicle deceleration in the target lane can not exceed a given safe limit b_{safe}. Risk criterion prevents critical lane changes and collisions, while the incentive criterion takes into account the advantages and disadvantages of other drivers associated with a lane change via the "politeness factor" p.

A lane change is shown in Figure 1. The MOBIL model depends on the two previous vehicles in the current and the target lanes, respectively. Thus, for a vehicle c considering a lane change, the previous vehicles in the target and current lanes are represented by n and o, respectively. The acceleration a_c denotes the acceleration of vehicle c on the current lane, and $\tilde{a}_c$ refers to the situation in the target lane, that is, to the new acceleration of vehicle c in the target lane. Likewise, $\tilde{a}_o$ and $\tilde{a}_n$ denote the acceleration of old and new previous vehicles after the lane change of vehicle c [10].

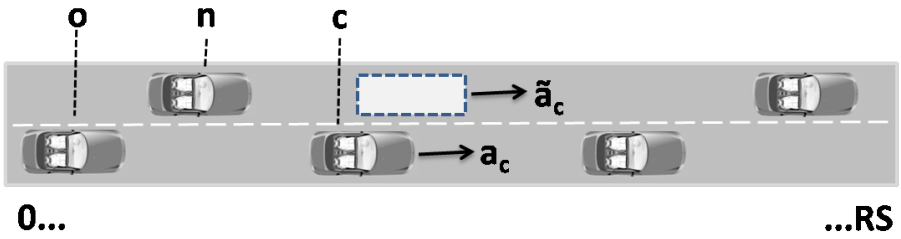

Fig. 1. Mobil notations. Adapted from [10].

According to [10], the incentive criterion determines if the lane change is better or not to a driver. In this model, the incentive is generalized to include the immediately affected neighbors. The politeness factor p determines to which degree these vehicles influence the lane- changing decision. Thus, the incentive criterion is given by the Equation 3.

$$\underbrace{\tilde{a}_c - a_c}_{\text{driver}} + p \left(\underbrace{\tilde{a}_n - a_n}_{\text{new behind}} + \underbrace{\tilde{a}_o - a_o}_{\text{old behind}} \right) > \Delta a_{th} \tag{3}$$

The first two terms of the Equation 3 denote the advantage of a possible lane change to the driver. The change is good if the driver can go faster in the new lane. The third term denotes the total advantage of the two immediately affected neighbors multiplied by the politeness factor p. The Δa_{th} term on the right-hand side represents a certain inertia and prevents lane changes if the overall advantage is only marginal compared with a "keep lane" directive.

2.3 Framework for Realistic Vehicular Mobility Models

For the purpose of guiding the developers through various challenges and options during the modeling, the authors of [13] propose a concept map for a comprehensible representation of a realistic vehicular mobility model. As can be seen in Figure 2, the concept map is organized around two major modules, **motion constraints** and the **traffic generator**. Additional modules such as **time** and **external influences** are also required for a fine tuning of the mobility patterns. The main modules (gray blocks) are implemented through several auxiliary modules (white ones), which are added according to the desired detail level. These last ones can be more explored in the original work. The main modules description are as follows [13]:

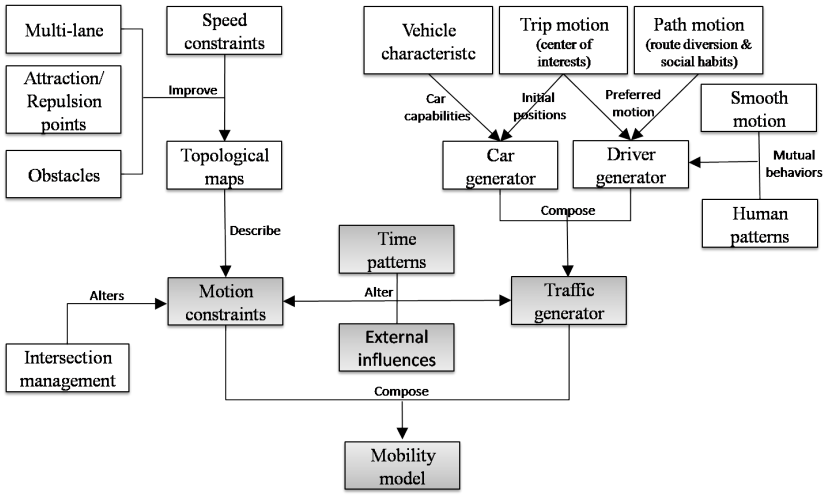

Fig. 2. Concept map of realistic mobility models. Adapted from [13].

- **Motion constraints** describe the relative degree of freedom available for each vehicle. Restrictions can be streets, buildings, vehicles and pedestrians.
- **Traffic generator** defines different kinds of vehicles, and handles their interactions according to the environment under study. Macroscopically, it models traffic densities, speeds and flows, while microscopically, it deals with properties such as the distance between cars, acceleration, braking, and overtaking.
- **Time** describes different mobility configurations for a specific time of the day. Traffic density is not uniform during a day. Peak times, such as rush hours or during special events, can be observed. This block influences the motion constraints and the traffic generator functional blocks.
- **External influences** model the impact of a communication protocol or any other source of information on the motion patterns. This block models the impact of accidents, temporary road works, or real-time knowledge of the traffic status on the motion constraints and the traffic generator blocks.

3 Probabilistic Model Checking

Probabilistic model checking is a formal, exhaustive and automatic technique for modeling and analyzing stochastic systems. PMC checks if the model satisfies a set of properties given in special types of logics.

A stochastic system M is usually a Markov chain or a Markov decision process. This means that the system must satisfy the Markov property, i.e., its behavior depends only on its current state and not on the whole system history, and each transition between states occurs in real-time.

Given a property ϕ expressed as a probabilistic temporal logic formula, PMC attempts to check whether a model of a stochastic system M satisfies the property ϕ with a probability $p \geq \theta$, for a probability threshold $\theta \in [0, 1]$.

Tools called model checkers such as PRISM [19] solve this problem. It requires two inputs: a modeling description of the system, which defines its behavior (for example, through the PRISM language), and a probabilistic temporal logic specification of a set of desired properties (ϕ).

The model checker builds a representation of the system M, usually as a graph-based data structure called Binary Decision Diagrams (BDDs), which can be used to represent boolean functions. States represent possible configurations, while transitions are changes from one configuration to another. Probabilities are assigned to the transitions between states, representing rates of negative exponential distributions.

Properties can be expressed quantitatively as "What is the shortest time which occurs overtaking?" or qualitatively as "Is overtake maneuver successful?", offering valuable insight over the system behavior.

Let $\mathbb{R}_{\geq 0}$ be the set of positive reals and AP be a fixed, finite set of atomic propositions used to label states with properties of interest. A labeled CTMC $\mathcal{C}$ is a tuple $(S, \bar{s}, \mathbf{R}, L)$ where:

- S is a finite set of states;
- $\bar{s} \in S$ is the initial state;
- $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is the transition rate matrix, which assigns rates between each pair of states;
- $L : S \to 2^{\mathrm{AP}}$ is a labeling function which labels each state $s \in S$ the set $L(s)$ of atomic propositions that are true in the state.

The probability of a transition between states s and s' being triggered within t time units is $1 - e^{-\mathbf{R}(s,s') \cdot t}$. The elapsed time in state s, before a transition occurs, is exponentially distributed with the *exit rate* given by $E(s) = \sum_{s' \in S} R(s, s')$. The probability of changing to state s' is given by $\frac{\mathbf{R}(s,s')}{E(s)}$.

Properties are specified using the Continuous Stochastic Logic (CSL) [23], which is based on the Computation Tree Logic (CTL) and the Probabilistic CTL (PCTL). The syntax of CSL formulas is the following:

$$\Phi \ ::= \ true \mid a \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathcal{P}_{\trianglelefteq p}[\phi] \mid \mathcal{S}_{\trianglelefteq p}[\phi]$$
$$\phi \ ::= \ \mathbf{X}\ \Phi \mid \Phi\ \mathbf{U}^{I}\ \Phi$$

where a is an atomic proposition, $\trianglelefteq \in \{>, <, \geq, \leq\}$, $p \in [0, 1]$ and $I \in \mathbb{R}_{\geq 0}$.

There are two types of CSL properties: transient ($\mathcal{P}_{\trianglelefteq p}$) and steady-state ($\mathcal{S}_{\trianglelefteq p}$). In this work we are interested in transient or time related properties. A formula $\mathcal{P}_{\trianglelefteq p}[\phi]$ states that the probability of the formula ϕ being satisfied from a state respects the bound $\trianglelefteq p$. Path formulas use the $\mathbf{X}$ (next) and the $\mathbf{U}^I$ (time-bounded until) operators. For example, formula $\mathbf{X}\,\Phi$ is true if Φ is satisfied in the next state.

This can be applied to check if a probability p is met for one property leading to other, such as $\mathcal{P}_{\trianglelefteq p}[\Phi_1 => \mathbf{X}\,\Phi_2]$, where Φ_1 and Φ_2 could be the properties "car reaches twice the truck's speed" and "*car* overtakes *truck* in 150 meters".

PRISM allows including **rewards** in the model, which are structures used to quantify states and transitions by associating real values to them. The state rewards are counted proportionately to the elapsed time in the state, while transition rewards are counted each time the transition occurs. In PRISM, rewards are described using the syntax:

rewards "*reward name*"
...
endrewards

Each reward is specified using the multiple reward commands syntax:

$$[sync]\, guard\, :\, reward\,;$$

Reward commands describe state and transition rewards. The *guard* predicate must be true. The *sync* is a label used to synchronize commands into a single transition. The *reward* is an expression that counts for the reward.

Reward properties can be used in states and transitions, e.g. "What is the expected reward (speed or throttle) for the car to travel 200 meters at time T?".

This reward can be instantaneous, obtaining its value at the given time through the property $\mathcal{R}_{=?}[\mathcal{I}^{=t}]$, or accumulated, calculating its value until the given time, using the property $\mathcal{R}_{=?}[\mathcal{C}^{<=t}]$. One can obtain the probability of a state reward by dividing it to the sum of all state rewards. The same procedure can be applied to transitions.

Rewards of paths in a Continuous-time Markov chain are summations of state rewards along the path and transition rewards for each transition between these states. State rewards are interpreted as the rate at which rewards are accumulated, essentially counting them, i.e. if t time units are spent in a state with state-reward r, the accumulated reward in that state is $r \times t$.

Another interesting PRISM feature, when reporting the result of model checking, is the ability to customize properties to obtain different results. This is done using filters, which use the following syntax:

$$filter(op, prop, states);$$

PRISM usually has to compute values for all states simultaneously, thus a specific point or all initial states can be selected. In the syntax, **op** is the filter

operator (e.g. max, min, avg), **prop** is any PRISM property and **states** is a Boolean-valued expression identifying a set of initial states to apply the filter.

4 Mobility VANET Model

Our model was created with a microscopic focus. The idea is to show the representation of movement of nodes through the analytical Equations 1, 2 and 3, previously described in Section 2.1. Signal propagation and communication have been abstracted. Our microscopic model take into account position, speed, and acceleration of the vehicles. For this, a overtaking vehicle scenario is implemented using the PRISM language. This has been done to demonstrate the viability of PMC usage to check microscopic aspects. Fragments of the models are presented below and the complete version can be found in the supplementary material and website [1].

Figure 3 illustrates the proposed scenario. There are three vehicles involved. The car $c1$ will overtake the truck, called **Leader**, which travels slower. However, the vehicle $c2$ is coming in the opposite direction. In this situation, $c1$ can not see $c2$, due to weather conditions or lack of attention. This scenario will happen in a 250 meters road. Thus, the model should answer questions such as "What is the probability of a collision?".

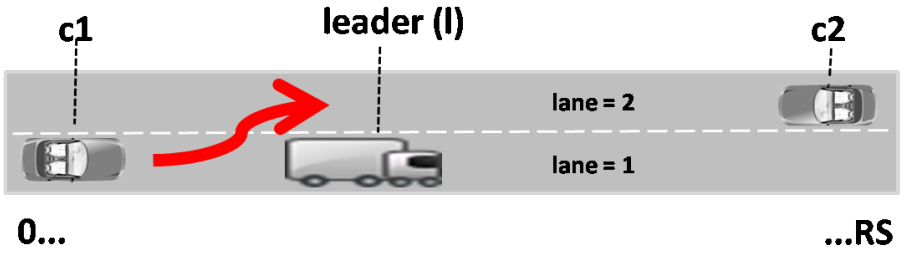

Fig. 3. Overtaking vehicle scenario.

The Figure 6 depicts the c1's variables (other vehicles are similar). Each vehicle maintains its current position and velocity. The variable *lane* informs where $c1$ is located. If the lane is equal to 1, then the vehicle is on right-hand side (default value), otherwise the car is on left-hand side. In other words, the vehicle is trying to overtake. The constants `desired_speed_car`, `desired_speed_truck` and `RS` (road side) constrain the model and they are respectively represented in m/s, m/s and m. The `carCrash` variable indicates whether $c1$ and $c2$ collided at some point in time.

An interesting feature of the model is that it does not have a specific initial state. This is achieved by the code shown in Figure 4. The restriction implemented states that vehicles $c1$ and $c2$ in opposite directions are separated by `RS` meters and that there is a leader (truck) between them, which will be overtaken by $c1$. However, the leader position and the initial speed of all involved can be

a combination of values. This creates several scenarios to be automatically explored. An interesting abstraction was adopted to c2's position. It starts in *one*, however, its real location on the road is given by $RS - pos\_c2$.

Initialization of Variables

```
init
  (pos_l >= truck_size + min_gap_car) & (pos_c1 = 1) & (pos_c2=1) &
  (lane = ( (v_c1>pos_l) | (a_c1 <= 0) ? 2 : 1) ) &
  (v_c1 >= 0 & v_c1 <= desired_speed_car) &
  (v_c2 >= 0 & v_c2 <= desired_speed_car) &
  (v_l >=0 & v_l <= desired_speed_truck) &
  (carCrash = false)
endinit
```

Fig. 4. Initial states for the model

The vehicles position is given by $x = x_i + vt + (a/2)t^2$, implemented in the PRISM language, which involves the initial position x_i, velocity vt, acceleration a, and time t. Each transition of the model represents a time period that is defined by the constant t. The acceleration of the vehicles are calculated by the IDM model previously presented in Section 2.1. The new speed is given by $v = v_i + at$ and it also depends on the vehicle acceleration. The Figure 5 describes a fragment of the model responsible for calculating the acceleration and position of vehicle $c1$. The formulas are similar for other vehicles.

As mentioned in Section 2.1, the IDM expression combines the free-road acceleration strategy, given by $a_{free}(v) = a[1 - (v/v_0)^{\delta}]$, with a deceleration strategy, given by $a_{brake}(s, v, \Delta v) = -a(s^*/s)^2$. Therefore, the Equation 1 has been algebraically split during implementation, because the vehicles do not suffer deceleration when there are no obstacles ahead. Thus, when the vehicle $c1$ overtakes the leader, $c1$ does not suffer slowdown, while the truck's acceleration, which used to have free way, starts to be influenced by the new $c1$'s position.

Acceleration and Position Formulas

```
formula a_c1_free = AM_car - AM_car * pow(v_c1 / desired_speed_car, exponent);
formula a_c1_obst = a_c1_free - a_brake_c1;
formula a_c1 = (overtook|lane=2?a_c1_free: (pos_l>=RS?a_c1_free:a_c1_obst));

formula a_brake_c1 = AM_car * pow(des_dyn_dis_c1 / deltaD_c1, 2);
formula des_dyn_dis_c1 = min_gap_car + max(0.0, v_c1 * T_car + (v_c1 * deltaV_c1) /
                                                              (2*pow(AM_car*BM_car,0.5) ));
formula deltaV_c1 = v_c1 - v_l;
formula deltaD_c1 = max(pos_l - pos_c1 - truck_size,1);//"max 1" to avoid division by zero

formula muv_c1 = (v_c1 + ( a_c1*pow(time,2)) / 2) > 0 ?
                 (v_c1 + (a_c1*pow(time,2)) / 2) : (-1 * (v_c1 + (a_c1*pow(time,2)) / 2));
```

Fig. 5. IDM model implementation

PRISM model comprises a set of modules which represent different components. The behavior of a module, i.e. the changes to its state that can occur, is specified by a set of *guarded commands*. These take the form:

$$[sync]guard \rightarrow rate : update;$$

where act is an (optional) action label, guard is a predicate over the variables of the model, rate is a (non-negative) real-valued expression and update is of the form:

$$(x'_1 = u_1)\&(x'_2 = u_2)\&\ldots\&(x'_k = u_k)$$

where $x_1; x_2; \ldots; x_k$ are local variables of the module and $u_1; u_2; \ldots; u_k$ are expressions over all variables.

Intuitively, a command is enabled in a global state of the PRISM model if the state satisfies the predicate guard. If a command is enabled, a transition that updates the module's variables according to update can occur with rate *rate*.

The modules `Mod_vC1` and `Mod_dC1` presented in Figure 6 are responsible for the transitions in the model which assign a new position and speed to vehicle $c1$, and also control the lane change of $c1$. If the vehicle is able to overtake according to the conditions presented by MOBIL model (refer to Subsection 2.2), the vehicle change to the left lane. If $c1$ is on the left lane and already overtook the leader, then $c1$ returns to the default lane. These modules are synchronized by label "m", which is placed inside the square brackets. The `Mod_dC1` is also responsible for detecting a crash, which happens when $c1$ and $c2$ are in the same lane and their coordinates are overlaid or the deceleration calculated by the Torricelle equation ($v^2 = v_i^2 + 2a\Delta x$) is unfeasible to be executed in a normal situation. The modules for the other vehicles involved are similar, although simpler because they just move forwards without overtake maneuvers.

Modules proposed

```
module Mod_vC1
  v_c1 : [0..desired_speed_car];  // speed

  [m] (pos_c1 <= RS) & (v_c1 <= desired_speed_car) ->
                          (v_c1' = min(max(ceil(v_c1 + a_c1)*time,0),desired_speed_car));
endmodule

module Mod_dC1
  pos_c1 : [1..RS];              // position
  lane     : [1..2]2;            //lane's c1 (1 - right lane, 2 - left lane)
  carCrash  : bool;

  [m] (pos_c1 <= RS) -> (pos_c1' = min( (ceil(pos_c1 + muv_c1)),RS) ) &
                        (lane' = ((lane = 2)&(pos_c1 >= (pos_l+min_gap_car+car_size)))?1:
                                                ((lane = 1)&(can_change_lane))?2:lane) &
                      (carCrash'=  ((CanotDecelaration | OverlapPosition) &
                                                (lane=2) & (carCrash=false) ) ?true:false);
endmodule
```

Fig. 6. Modules implementation

5 Results

Finally, the model built using the PRISM language can be verified. The idea is to check the correctness of IDM code and analyze different situations about the modeled scenario. The experiments have been performed in an Intel(R) Xeon(R) CPU X3323 , 2.50 GHz which has 16 GB of RAM memory. The model presented has 386 243 states, 386 243 transitions and 38 400 initial states. For some properties we have varied the number of initial states through *filters*. The longest

time to build the model was 2 360.838 s. The longest time to check a property was for Property 8 of Figure 11, taking 5.418 s.

In order to analyze some situations about the scenario, several interesting questions can be made. For example, the first property (Figure 7) checks the probability of a car-crash. The result was: `[0.0, 1.0]` for a range of values over initial states. The answer shows that there are situations without accident, however there are cases of car-crash.

The third property (Figure 7) checks the average probability of an accident taking into account all initial states. Thus, this scenario has a 98% chance of collision. The fourth property only confirms the results of these two previously mentioned properties. It is a non-probabilistic query and the result was *true* for the question "Are there situations without accidents?". The **E** (Exists) operator asks whether some path from a state satisfies a particular path formula. If the result is true, a witness will be generated. In this case, it was provided the following counter-example: `(0, 0, 0, 1, 1, false, 1, 19)`, which represents the initial state with values for the respective variables `v_c1`, `v_c2`, `v_l`, `pos_c1`, `lane`, `carCrash`, `pos_c2` and `pos_l`.

The second property shows another analysis, having calculated the result `[0.0,1.0]` considering the range of values over initial states for the question "Is it possible to finish the scenario without overtake?", in other words, the leader reaches the finish before `c1`. Thus, there are cases with and without overtake.

Car-crash Scenario Properties

```
(1) P=? [ F (carCrash=true) ]
What is the probability of an accident occurs?

(2) P=? [ ((pos_c1<RS & carCrash=false) U (pos_l>=RS & carCrash=false))]
What is the probability of not occurring overtakes in this scenario?

(3) filter(avg, P=? [ F carCrash=true ], "init")
What is the average probability of an accident occurs?

(4) E [ F (carCrash=false) ]
Is there, at least, one path which does not lead to the accident?
```

Fig. 7. Properties of Overtake Maneuver

As we have mentioned above, the operator **E** generates a counter-example (a path reaching the "goal" state). Using this witness, for instance for Property 4, we can analyze in detail the situation of accidents in the scenario. Since we have included rewards in our model, we are able to quantify the speed, acceleration and movement over time using the **I** (instant) operator. Some implemented rewards and properties are shown in Figure 8, the latter using the *filter* command to check specifically the counterexample available. The operator R is the responsible to get the reward values.

Figure 9 shows the result of analysis, showing the position of the three vehicles over time. The red line varies between 10 and 20 and it represents the lane of the vehicle *car1* during overtaking. The first value means that *car1* is in the default lane (right lane), the value 20 means that the vehicle is traveling in the left lane to overtake. The Figure 9.1 shows the behavior of vehicles without collision. Note

Movement and Lane Rewards

```
rewards "dLeader"
  true : pos_l;
endrewards

rewards "dCar"
  true : pos_c1;
endrewards

rewards "dCarOpposite"
  true : pos_c2;
endrewards

rewards "laneCar"
  true : lane;
endrewards
```

Movement and Lane Quantitative Properties

```
(5) filter(max, R{"dCar"}=? [ I=T ],
      (pos_l=19)&(pos_c1=1)&(pos_c2=1)&(lane=1)&
      (v_c1=0)&(v_c2=0)&(v_l=0)&(carCrash=false))

What is the expected distance reward for the vehicle c1 on
the road at time T?

(6) filter(max,R{"laneCar"}=? [ I=T ],
        (pos_l=19)&(pos_c1=1)&(pos_c2=1)&(lane=1)&
        (v_c1=0)&(v_c2=0)&(v_l=0)&(carCrash=false))

What is the expected lane reward for the vehicle c1 on the
road at time T?

(7) filter(max,RS-R{"dCarOpposite"}=? [I=T],
        (pos_l=50)&(pos_c1=1)&(pos_c2=1)\&(lane=1)&
        (v_c1=0)\&(v_c2=0)\&(v_l=0)&(carCrash=false))

What is the expected distance reward for the vehicle c2 on
the road at time T? (With different initial conditions)
```

Fig. 8. Movement and Lane Rewards, and Quantitative Properties.

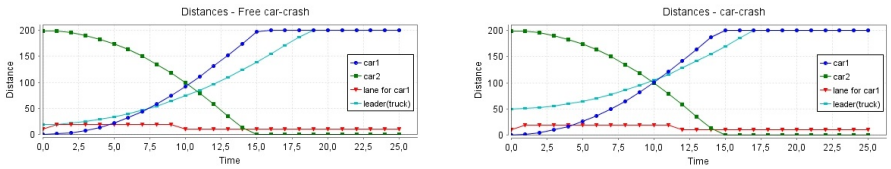

9.1 Motion in a normal overtake 9.2 Motion with accident

Fig. 9. Scenario analysis

that *car1* overcomes the leader at the instant 7.5 and when the positions of *car1* and *car2* overlap, the first car already returned to the right lane. However, in the Figure 9.2 can be seen that the positions overlap at time 10 and *car1* is in the left lane, meaning that there was a collision.

Figure 10 shows the evolution of the acceleration and velocity of *car1* and *leader* (truck) in the scenario of overtaking without collision. Speeds rise according to acceleration until reaching the maximum limit of the road. As the acceleration and the speed limit are lower, the car can overtake easier. It is interesting to note that the acceleration modeled with IDM is affected by lane change of *car1*. Right at the instant 2, the acceleration of the *car1* rises abruptly, because in this moment, the driver concludes to be more advantageous changing to the left lane, instead of maintaining in its lane. As the *car1* is reaching the desired speed, the acceleration is decreasing, which happens linearly. The truck also reduces the acceleration linearly as the desired speed is reached. At time 8, the deceleration is slightly more accentuated due to the entrance of the *car1* on the default lane, as soon as the overtaking is completed.

Analysis regarding to the time spent during overtake or going through the entire route for each vehicle can also be computed. The Figure 11 shows two examples of this type of verification. These properties use the reward "step", responsible for providing the value 1 for each change of state in the model, which is equivalent to 1 second in a real scenario. These properties use the operator

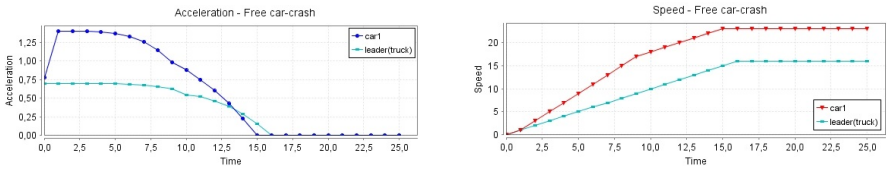

10.1 Acceleration evolution 10.2 Speed evolution

Fig. 10. Analysis in a free car-crash overtake

F (reachability), which is associated with the reward "step". According to [18], the reward property "**F** prop" corresponds to the reward accumulated along a path until a satisfactory state is reached. In the case, where the probability of reaching a state satisfying prop is less than 1, the reward is equal to infinity.

Time Reward

```
rewards "steps"
  true : 1;
endrewards
```

Time Properties

```
(8) filter(print, R{"steps"}=? [ F carCrash=true ], 'init')
The lowest and highest time to c1 collide during its journey.

(9) filter(print, R{"steps"}=? [ F carCrash=false & pos_c1=RS
], 'init')
The maximum and minimum time to c1 completes the path without
collision.
```

Fig. 11. Time Rewards and Properties

Property 8 calculates the overtake time, which results in a possible collision for all initial states, thus the presented result was a value range of `[8.0, infinity]` seconds. The infinity value represents the initial states without collision, i.e. initial states that have probability less than 1. Therefore, to find the maximum time limit for the collision it is enough to analyze the PRISM output log file which will have the travel time for all initial states and their successors, which is available due to the parameter "print" in the *filter* command.

The range of minimum and maximum time of collision is `[8.0, 11.0]`, i.e. the shortest time of an accident is 8 seconds and the greatest time is at instant 11. Thus, they can be simulated, respectively with the following initial states `(1, 6, 1, 2, 2, false, 6, 20)` and `(0, 0, 1, 1, 1, false, 1, 19)`, for the following variable `v_c1`, `v_c2`, `v_l`, `pos_c1`, `lane`, `carCrash`, `pos_c2` and `pos_l`.

In a similar way, the Property 9 calculates the minimum and maximum time of a successful overtake. This also takes into account all possible initial states. Thus, the range of values presented were between `[13.0, Infinity]` seconds. Again, for all initial states which the probability of **F** to be satisfied is less than 1, it is assigned the infinity value. Thus, analyzing the PRISM log file, we can identify the new range of values, and the infinity value is `[13.0, 16.00]`. Their respectively counter-examples are `(4, 0, 0, 1, 1, false, 1, 40)` and `(0, 0, 0, 1, 1, false, 1, 19)`, for the same variables presented in Property 8.

6 Conclusions

It is essential to test and analyze VANETs in order to prevent loss of life. Simulations are used to check protocols and applications, however, they have to deal

with two unconnected worlds – network and traffic – which must work together. In this context, there are challenges that must be addressed by the academic community. A complementary tool to simulations is model checking, a technique that automatically and exhaustively explores a model. However, researchers can use simulation to large-scale analysis and model checking to test thoroughly in a smaller proportion. Thus, they can supply solutions to known problems for simulations and model checking, such as determining exact probabilities and avoiding the state explosion, respectively.

In this article we have presented the formal modeling and analysis of mobility models using probabilistic model checking to represent an overtake situation. A microscopic vision was presented to provide a detailed analysis. This was possible using analytical formulas to represent position, speed, and acceleration. The model shows that there is a huge chance of an accident (98% in some scenarios), however there are situations without collision.

In general, during implementation we have noticed some limitations in the PRISM language, e.g., the absence of some mathematical functions, the lack of subroutine (function and procedure) and formal parameters. This fact impairs the legibility of the model and makes difficult to implement and maintain the models. However, the IDM and MOBIL models can be perfectly implemented and used in PRISM.

The implementation of motion provides important information such as instantaneous speed, acceleration and position through rewards, besides answering questions regarding the probability of events. Our model follows the framework shown in Figure 2, presenting smooth motion and human driving patterns, furthermore following speed constraints and considering obstacles. All of this is provided by IDM and the MOBIL.

The motion patterns are not considered because we are analyzing specific situations instead of a large flow of vehicles. Furthermore, the mobility modules can be easily coupled with network protocols. In addition, the modeling is easily adaptable under various situations, such as multilane highway or an intersection. For example, to implement a curve road with a higher abstraction level, it is simply necessary to change the limited speed to a value less than a straight road, thus vehicles will reduce the speed while they are crossing a curve.

Future works: explore more mobility scenarios following the concepts and examples presented here and couple them with models that represent communication and signal propagation using a probabilistic method, such as [6], thus, making it possible to do a complete analysis of the VANET in a stochastic way.

References

1. `http://www.dcc.ufmg.br/~bruno.ferreira/sbmf2014/`
2. Alves, R., et al.: Redes veiculares: Princípios, aplicações e desafios. in Minicursos do Simpósio Brasileiro de Redes de Computadores - SBRC'2009 (May 2009)
3. Bai, F., Krishnan, H., Sadekar, V., Holland, G., Elbatt, T.: Towards Characterizing and Classifying Communication-based Automotive App. from a Wireless Networking Perspective. In: In Proc. of IEEE Workshop on Automotive Networking (2006)

4. Boban, M., Vinhoza, T.T.V.: Modeling and simulation of vehicular networks: Towards realistic and efficient models. In: Mobile Ad-Hoc Networks: Applications. Intech (2011)
5. Bouassida, M.S., Shawky, M.: A cooperative congestion control approach within vanets: formal verification and performance evaluation. EURASIP J. Wirel. Commun. Netw. 2010, 11:1–11:12 (apr 2010)
6. Boulis, A., Fehnker, A., Fruth, M., McIver, A.: Cavi – simulation and model checking for wireless sensor networks. QEST (2008)
7. Choffnes, D.R., Bustamante, F.E.: An integrated mobility and traffic model for vehicular wireless networks. In: Laberteaux, K.P., Hartenstein, H., Johnson, D.B., Sengupta, R. (eds.) Vehicular Ad Hoc Networks. pp. 69–78. ACM (2005)
8. Christian, A.: Reliable model checking for wsns. In: Proc. of the 8th GI/ITG KuVS Fachgesprach (2009)
9. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press (1999)
10. Dirk Helbing, T., Martin, A.K.: General lane-changing model mobil for car-following models. In: Transp. Research Record: Journal of the Transp. Research Board. pp. 86–94. Transp. Research Board of the National Academies (2007)
11. Ferreira, M., Fernandes, R., Conceição, H., Viriyasitavat, W., Tonguz, O.K.: Self-organized traffic control. VANET '10, ACM, New York, NY, USA (2010)
12. Gipps, P.G.: A model for the structure of lane-changing decisions. Transportation Research Part B: Methodological 20(5), 403–414 (1986)
13. Harri, J., Filali, F., Bonnet, C.: Mobility models for vehicular ad hoc networks: a survey and taxonomy. IEEE Communications Surveys & Tutorials 11 (Dec 2009)
14. Hartenstein, H., Laberteaux, K., Ebrary, I.: VANET: vehicular applications and inter-networking technologies. Wiley Online Library (2010)
15. Helbing, D.: Traffic and related self-driven many-particle systems. Rev. Mod. Phys. 73, 1067–1141 (Dec 2001)
16. Kesting, A., Treiber, M., Helbing, D.: Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. Royal Society of London Philosophical Transactions Series A 368, 4585–4605 (Sep 2010)
17. Khosroshahy, M.: IEEE 802.11 and propagation modeling: A survey and a practical design approach (2007)
18. Kwiatkowska, M., Norman, G.: PRISM - Property Specification. `http://www.prismmodelchecker.org/manual/PropertySpecification/Reward-basedProperties` (2011), access date: 28 jan. 2014
19. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Proc. CAV. Springer (2011)
20. Lomuscio, A., Strulo, B., Walker, N.G., Wu, P.: Model checking optimisation based congestion control algorithms. Fundam. Inf. 102(1) (Jan 2010)
21. M. Althoff, O. Stursberg, M.B.: Safety assessment of driving behavior in multi-lane traffic for autonomous vehicles. In: Proc. of the IEEE Intelligent Vehicles Symposium. Shaanxi, China (June 2009)
22. Mangharam, R., Weller, D.S., Stancil, D.D., Rajkumar, R., Parikh, J.S.: Groovesim: a topography-accurate simulator for geographic routing in vehicular networks. In: Vehicular Ad Hoc Networks. pp. 59–68. ACM (2005)
23. Parker, D.: Implementation of Symbolic Model Checking for Probabilistic Systems. Ph.D. thesis, University of Birmingham (2002)
24. Treiber, M., Hennecke, A., Helbing, D.: Congested traffic states in empirical observations and microscopic simulations. Rev. E 62, Issue 62, 2000 (2000)

Towards completeness in Bounded Model Checking through Automatic Recursion Depth Detection

Grigory Fedyukovich and Natasha Sharygina

Faculty of Informatics, University of Lugano
Via Guiseppe Buffi 13, CH-6904 Lugano, Switzerland

Abstract The presence of recursive function calls is a well-known bottleneck in software model checking as they might cause infinite loops and make verification infeasible. This paper proposes a new technique for sound and complete Bounded Model Checking based on detecting depths for all recursive function calls in a program. The algorithm of detection of recursion depth uses over-approximations of function calls. It proceeds in an iterative manner by refining the function over-approximations until the recursion depth is detected or it becomes clear that the recursion depth detection is infeasible. We prove that if the algorithm terminates then it guarantees to detect a recursion depth required for complete program verification. The key advantage of the proposed algorithm is that it is suitable for generation and/or substitution of function summaries by means of Craig Interpolation helpful to speed up consequent verification runs. We implemented the algorithm for automatic detection of recursion depth on the top of our SAT-based model checker `FunFrog` and demonstrate its benefits on a number of recursive C programs.

1 Introduction

Model checking plays an important role in both proving program correctness and finding bugs. It provides a powerful fully automated engine which is able to search for an assertion violation among all possible combinations of the input values. These advantages are however hindered by the high complexity of analysis, known as the state-space explosion phenomenon. To combat this problem, many effective state-space reduction solutions have been developed to allow model checking to scale to verification of complex systems. The most successful solutions are symbolic model checking among which are Bounded Model Checking (BMC) [2], and abstraction-based approaches such as predicate abstraction [8], interpolation-based reasoning [11], and function summarization [12,13,1,19].

BMC has been shown to be particularly successful in safety analysis of software. The state-of-the-art BMC-based tools such as `CBMC` [3], `LLBMC` [14], `VeriSoft` [9], `FunFrog` [18], just to name a few, have been successfully applied to verification of industrial-size programs. The well-known limitation of BMC is that it is aimed at searching for errors in a program within the given number

(bound) of loop iterations and recursion depth. For this reason, BMC is suitable only for program falsification, while for complete verification it requires finding a sufficient bound. This problem remains open: the BMC tools analyze an under-approximation of a program using some particular bound, defined a priori by the user or set by the tool to some constant, and check the program only up to this bound.

There exists a number of (direct and indirect) solutions for the automatic loop bound detection (i.e., constant propagation, k-induction, loop summarization, etc). However, dealing with recursive function calls is more complicated and more expensive in practice. This paper proposes an approach for the automatic recursion depth detection in BMC and shows its applicability in practice.

In particular, we present a BMC algorithm enhanced with automated construction of the sufficient unwinding[1]. The algorithm iteratively explores the program calltree and over-approximates recursive function calls while treating precisely the other ones. The entire abstraction of the calltree is then checked on-the-fly with respect to a given assertion. If the assertion holds in the current level of abstraction then the corresponding unwinding is sufficient to guarantee complete verification (and the length of the longest unwinding chain constitutes the recursion depth). Otherwise, the algorithm identifies which over-approximated function calls are responsible for the assertion violation. These function calls are going to be refined and the algorithm goes to the next iteration.

Our approach is developed to reach efficiency in BMC. At each iteration, it refines only a minimal set of over-approximated function calls, i.e., only those responsible for spuriousness of the error on the previous iteration. Clearly, the algorithm is not guaranteed to terminate when there are unbounded sequences of recursive calls in the program. But if for every possible value of input parameters, every recursive function in the real program is called a fixed number of times, the algorithm automatically detects this number and terminates.

We further demonstrate how our algorithm can be made practical by extending our earlier work on construction and reusing of interpolation-based function summaries in BMC [19] for checking different assertions. In the current work, aside from checking user-provided assertions, we use a heuristic called *assertion decomposition* to artificially implant *helper*-assertions into the recursive program. These assertions are then checked incrementally to generate function summaries that will be reused to speed up verification of the user-provided assertions.

We implemented the approach on the top of `FunFrog` BMC, previously restricted to work only for a user-supplied recursion depth. We evaluated it on a range of academic and industrial recursive programs requiring bitwise and non-linear reasoning. Our experimentation confirmed that the summarization-based recursion depth detection in many cases makes BMC complete and dramatically improves its performance compare to the classical BMC approach (e.g., `CBMC`).

Algorithmically, the closest body of work is the `Corral` [10] tool (see related work section for detailed comparison). It is a solver for a restricted version of the

[1]The algorithm relies on the output of a loop bound detection routine (e.g., conversion loops to recursion) done by an external tool or set by the user.

reachability-modulo-theories problem, and it also uses summaries in its bounded analysis to guarantee a practical solution. Unlike in our approach, in the `Corral`, 1) the depth of recursion is bounded by a user-supplied recursion depth and 2) an external tool [7] is used to generate function summaries which in general may not be helpful to verify the given assertion. Our approach is able to generate relevant function summaries by itself. Moreover, it forces summaries to be bit-precise and highly related to the given assertion. It makes our algorithm converge more effectively and faster.

The rest of the paper is structured as follows. Sect. 2 defines the notation and presents background on BMC, function summarization and refinement. Sect. 3 presents the BMC algorithm with automatic detection of recursion depth, proves its correctness and demonstrates its application to function summarization-based model checking. Sect. 4 discusses different experimentation scenarios of the approach including the assertion decomposition heuristic. Sect. 5 provides a comparison with the related work and Sect. 6 concludes the paper.

2 Preliminaries and Previous Work

We first define basic constructs required to present the new algorithm. In particular, we explicitly define recursion, function summaries and basic BMC steps.

2.1 Programs, function calls, recursion depth

Definition 1 (cf. [19]). *An* unwound program *for a depth ν is a tuple $P_\nu = (\hat{F}_\nu, \hat{f}_{main}, child)$, such that $\hat{F}_\nu$ is a finite set of function calls, unwound up to the depth ν, $\hat{f}_{main} \in \hat{F}_\nu$ is a program entry point and $child \subseteq \hat{F}_\nu \times \hat{F}_\nu$ relates each function call $\hat{f}$ to all function calls invoked directly from it.*

There is a fixed set F to represent functions declared in the program and a possibly unbounded set $\hat{F}$ to represent function calls. A call $\hat{f} \in \hat{F}$ corresponds to a call of a target function, determined by a mapping $target : \hat{F} \to F$. A subset $\hat{F}_\nu \subseteq \hat{F}$ is introduced to help handling recursion. There is exactly one call of function f_{main}, but there may be several calls of the other functions. For simplicity, later we will use primes (i.e., $\hat{f}'$, $\hat{f}''$,..) and indexes (i.e., $\hat{f}_1$, $\hat{f}_2$,..) to differentiate the calls of the same function $f \in F$ in the unwound program.

The set of function calls $\hat{F}$ together with the relation *child* can be represented by a corresponding calltree with the root $\hat{f}_{main}$. We also use relation $subtree \subseteq \hat{F} \times \hat{F}$, a reflexive transitive closure of *child*. Now we can define recursive functions using this notation.

Definition 2. *A function f is* recursive *if for every call $\hat{f}_i$, there is another call $\hat{f}'_i$ in its subtree, and $target(\hat{f}_i) = target(\hat{f}'_i) = f$.*

According to Def. 2, the calltree of a program with recursive functions is infinite. As detailized later in this section, for classical BMC it has to be bounded. A recursive function f is unwound ν times if there is a sequence of function calls

(later called an *unwinding chain*) $\hat{f}_0$, $\hat{f}_1$,.. $\hat{f}_\nu$, where $1 \leq i \leq \nu$, $target(\hat{f}_i) = f$, and each $\hat{f}_{i+1}$ is in the subtree of $\hat{f}_i$. The set of function calls $\hat{F}_\nu$ and the relation *child* define a finite corresponding calltree. If there are no recursive function calls in the program $P_\nu = (\hat{F}, \hat{f}_{main}, child)$ then $\hat{F}_\nu \equiv \hat{F}$ for any ν.

BMC is aimed at checking assertions in a program within the given bound of loop iterations and recursion depth. If the unwinding number ν is provided a priori, BMC unrolls the loops and recursion up to ν, encodes the program symbolically and delegates the checking to a SAT solver. If the number is not provided a priori, BMC may go into an infinite loop and not terminate. Typically in the absence of the number or when the number is set too high, a predefined timeout is used to cope with this problem.

BMC encodes the program into the Static Single Assignment (SSA) form, where each variable is assigned at most once. The SSA form is then conjoined with the negation of the assertion condition and converted into a logical formula, called a *BMC formula*. The BMC formula is checked for satisfiability, and every its satisfying assignment identifies an error trace. Otherwise, the program is safe up to ν. Notably, this unwinding number may not be sufficient for complete verification. A program can be proven safe for ν, but buggy for $\nu + 1$.

Fig. 1 illustrates BMC encoding for a simple C program (Fig. 1a) with a recursive function `f`. For this example, the recursion depth $\nu = 5$ guarantees complete verification.[2] In this setting, it is assumed that this recursion depth is

```
int f(int a) {
  if (a < 10)
    return f (a + 1);
  return a - 10;
}

void main() {
  int y = 1;
  int x = nondet();

  if (x > 5)
    y = f(x);

  assert(y >= 0);
}
```

(a) C code

```
y0 = 1;
x0 = nondet();
if (x0 > 5) {
  a0 = x0;
  // f (unwind 1)
  if (a0 < 10)
    // f (unwind 2)
    ...
    // end f (unwind 2)
    ret0 = ...;
  else
    ret1 = a0 - 10;
  ret2 = phi(ret0, ret1);
  // end f (unwind 1)
  y1 = ret2;
}
y2 = phi(y0, y1);
assert(y2 >= 0);
```

(b) SSA form

$$
\begin{aligned}
&y_0 = 1 \wedge \\
&x_0 = nondet_0 \wedge \\
&a_0 = x_0 \wedge \\
&ret_0 = ... \wedge \\
&... \wedge \\
&ret_1 = a_0 - 10 \wedge \\
&(x_0 > 5 \wedge a_0 < 10 \Rightarrow \\
&\qquad ret_2 = ret_0) \wedge \\
&(x_0 > 5 \wedge a_0 \geq 10 \Rightarrow \\
&\qquad ret_2 = ret_1) \wedge \\
&y_1 = ret_2 \wedge \\
&(x_0 > 5 \Rightarrow y_2 = y_1) \wedge \\
&(x_0 \leq 5 \Rightarrow y_2 = y_0) \wedge \\
&y_2 < 0
\end{aligned}
$$

(c) BMC formula

Figure 1: BMC formula generation

[2] see more details on termination in Sect. 3.1

given a priori. During unwinding (Fig. 1b), a call of function `f` is substituted by its body. There will be five such nested substitutions, and the sixth call is simply skipped in the example. The encoded BMC formula is shown on Fig. 1c.

Classical BMC algorithms use a monolithic BMC formula, as described in details in [3]. For specialized BMC algorithms (such as in our earlier work on function summarization [19] and upgrade checking [6], and the new algorithm for automatic detection of recursion depth) it is convenient to use a so called *Partitioned* BMC formula, which is going to be presented in Sect. 2.2.

2.2 PBMC encoding

Definition 3 (cf. [19]). *Let $\hat{F}_\nu$ be an unwound calltree, π encodes an assertion, $\phi_{\hat{f}}$ symbolically represent the body of a function f, a target of the call $\hat{f}$. Then a* partitioned BMC *(PBMC) formula is constructed as $\neg\pi \wedge \bigwedge_{\hat{f}\in\hat{F}_\nu} \phi_{\hat{f}}$.*

Fig. 2 demonstrates creation of a PBMC formula for the example from Fig. 1a. In the example program, unwound 5 times, the partitions for function calls $\mathtt{f}_1$,$\mathtt{f}_2$,..$\mathtt{f}_5$ and `main` are generated separately. They are bound together using a special boolean variable $callstart_{\hat{f}}$ for every function call $\hat{f}$. Intuitively, $callstart_{\hat{f}}$ is equal to *true* iff the corresponding function call $\hat{f}$ is reached. Note that the assertion π is not encoded inside $\phi_{\hat{f}_{main}}$, as in classical BMC, but separated from the rest of the formula, such that it helps interpolation.[3]

Formula $\phi_{\hat{f}_1}$ that encodes the function call $\mathtt{f}_1$ aims to symbolically represent the function output argument ret_0 by means of the function input argument a_0, symbolically evaluated in $\phi_{\hat{f}_{main}}$. At the same time, $\phi_{\hat{f}_1}$ relies on the value of ret_3 defined in $\phi_{\hat{f}_2}$ by means of a_1. Similar reasoning is applied to create each of the following partitions: $\phi_{\hat{f}_2}$,.. $\phi_{\hat{f}_5}$.

$$y_0 = 1 \wedge$$
$$x_0 = nondet_0 \wedge$$
$$a_0 = x_0 \wedge$$
$$x_0 > 5 \Leftrightarrow callstart_{\hat{f}_1} \wedge$$
$$y_1 = ret_0 \wedge$$
$$(x_0 > 5 \Rightarrow y_2 = y_1) \wedge$$
$$(x_0 \leq 5 \Rightarrow y_2 = y_0)$$

(a) formula $\phi_{\hat{f}_{main}}$

$$y_2 \geq 0 \Leftrightarrow \pi$$

(b) definition of π

$$(a_0 < 10 \Leftrightarrow callstart_{\hat{f}_2}) \wedge$$
$$a_1 = a_0 + 1 \wedge$$
$$ret_1 = ret_3 \wedge$$
$$ret_2 = a_0 - 10 \wedge$$
$$(callstart_{\hat{f}_1} \wedge a_0 < 10 \Rightarrow ret_0 = ret_1) \wedge$$
$$(callstart_{\hat{f}_1} \wedge a_0 \geq 10 \Rightarrow ret_0 = ret_2)$$

(c) formula $\phi_{\hat{f}_1}$

Figure 2: PBMC formula generation

[3] see more details on interpolation in Sect. 2.3

2.3 Craig Interpolation and Function Summarization

Definition 4 (cf. [4]). *Given formulas A and B, such that $A \wedge B$ is unsatisfiable.* Craig Interpolant *of A and B is a formula I such that $A \rightarrow I$, $I \wedge B$ is unsatisfiable and I is defined over the common alphabet to A and B.*

For mutually unsatisfiable formulas A and B, an interpolant always exists [4]. For quantifier free propositional logic, an interpolant can be constructed from a proof of unsatisfiability [16]. Interpolation is used to generate function summaries to speed up incremental verification (see our earlier work [19,18]).

Definition 5 (cf. [19]). Function summary *is an over-approximation of the function behavior defined as a relation over its input and output variables.*

A summary contains all behaviors of the function and (due to its over-approximating nature) possibly more. The infeasible behaviors (detected during analysis of abstract models) have to be refined by means of the automated procedure, as will be described in Sect. 2.4.

If the program is safe with respect to an assertion π, then the PBMC formula representing the program is unsatisfiable. The interpolation procedure is applied repeatedly for each function call $\hat{f}$. It splits the PBMC formula into two parts, $\phi_{\hat{f}}^{subtree}$ and $\phi_{\hat{f}}^{env}$ (1). The former encodes the subtree of $\hat{f}$. The latter corresponds to the rest of the encoded program including a negation of assertion π.

$$\phi_{\hat{f}}^{subtree} \equiv \bigwedge_{\hat{g} \in \hat{F}: subtree(\hat{f},\hat{g})} \phi_{\hat{g}} \qquad \phi_{\hat{f}}^{env} \equiv \neg\pi \wedge \bigwedge_{\hat{h} \in \hat{F}: \neg subtree(\hat{f},\hat{h})} \phi_{\hat{h}} \qquad (1)$$

Since $\phi_{\hat{f}}^{subtree} \wedge \phi_{\hat{f}}^{env}$ is unsatisfiable, the proof of unsatisfiability can be used to extract an interpolant $I_{\hat{f}}$ for $\phi_{\hat{f}}^{subtree}$ and $\phi_{\hat{f}}^{env}$. Such formula $I_{\hat{f}}$ is then considered as a summary for the function call $\hat{f}$. While verifying another assertion π', the entire part $\phi_{\hat{f}}^{subtree}$ of the PBMC formula will be replaced by the summary formula $I_{\hat{f}}$.

2.4 Counter-Example Guided Refinement

Definition 6 (cf. [19]). *A* substitution scenario *for function calls is a function $\Omega : \hat{F} \rightarrow \{inline, sum, havoc\}$.*

For each function call, a substitution scenario determines a level of approximation as one of the following three options: *inline* when it processes the whole function body; *sum* when it substitutes the call by an existing summary, and *havoc* when it treats the call as a nondeterministic function. Since *havoc* abstracts away the function call, it is equivalent to using a summary *true*.

In the incremental abstraction-driven analyses [19,6], substitution scenarios are defined recurrently. Algorithms start with the least accurate *initial* scenario

Ω_0, and iteratively *refine* it. In (2) and (3), we adapt the definitions from [19] to the recursive case.

$$\Omega_0(\hat{f}) = \begin{cases} \mathit{sum}, & \text{if there exists a summary of } \hat{f} \\ \mathit{inline}, & \text{if } \hat{f} \text{ is not recursive or } \nu \text{ is not exceeded} \\ \mathit{havoc}, & \text{if } \hat{f} \text{ is recursive and } \nu \text{ is exceeded} \end{cases} \tag{2}$$

$$\Omega_{i+1}(\hat{f}) = \begin{cases} \mathit{inline}, & \text{if } \Omega_i(\hat{f}) \neq \mathit{inline} \text{ and } \mathit{callstart}_{\hat{f}} = \mathit{true} \\ \Omega_i(\hat{f}), & \text{otherwise} \end{cases} \tag{3}$$

When a substitution scenario Ω_i leads to a satisfiable PBMC formula (i.e., there exists an error trace ϵ), an analysis of ϵ is required to shows that the error is either real or spurious. By construction of the PBMC formula, for each function call $\hat{f}$, a variable $\mathit{callstart}_{\hat{f}}$ is evaluated to *true* iff $\hat{f}$ appears along ϵ. Consequently, each $\hat{f}$ might be responsible for spuriousness of ϵ if $\hat{f}$ was not precisely encoded and $\mathit{callstart}_{\hat{f}} = \mathit{true}$. If there is no function call, satisfying the above mentioned conditions, ϵ is real and must be reported to the user.

3 Bounded Model Checking with Automated Detection of Recursion Depth

This section presents an iterative abstraction-refinement algorithm for BMC with automated detection of recursion depth. We first present a basic algorithm, where all function calls are treated nondeterministically (Sect. 3.1). Then we strengthen this algorithm to support generation and use of interpolation-based function summaries (Sect. 3.2).

3.1 Basic Algorithm

An overview of the algorithm is depicted in Alg. 1. The algorithm starts with a preset recursion depth ν[4] and iterates until it detects the actual recursion depth, needed for complete proof of the program correctness, or a predefined timeout is reached. Notably, at each iteration of the algorithm, ν gets updated and is equal to the length of the longest unwinding chain of recursive function calls. In the end of the algorithm, all recursive calls are unwound exactly same number of times as they would be called during the execution of the program.

The details of the computation are given below. First, the algorithm aims to construct a PBMC formula ϕ using the sets $\hat{F}_\nu$ and $\mathbb{T}$. Every function call $\hat{f} \in \hat{F}_\nu$ is encoded precisely, every function call $\hat{g} \in \mathbb{T}$ is treated nondeterministically. In particular, bodies of function calls from set $\hat{F}_\nu$ are encoded into the SSA forms (i.e., method `CreateFormula`) and put together into separate partitions (one partition per each function call) of ϕ (line 3). At the same time, all function

[4]The algorithm can be initialized with any number value as demonstrated in our experiments.

Algorithm 1: BMC with automatic detection of recursion depth

Input: Initial recursion depth ν; Program unwound ν times: $P_\nu = (\hat{F}_\nu, \hat{f}_{main}, child)$; Assertion to be checked: π; $TimeOut$
Output: Verification result: $\{SAFE,\ BUG,\ TimeOut\}$; Detected recursion depth ν; Error trace: ϵ
Data: ϕ: PBMC formula; $\mathbb{T}$: temporary set of function calls to be refined

```
1  while ¬TimeOut do
2      T ← {ĝ ∉ F̂_ν | child(f̂, ĝ), f̂ ∈ F̂_ν};                    // get refinement candidates
3      φ ← ¬π ∧ ⋀_{f̂ ∈ F̂_ν} CreateFormula(f̂) ∧ ⋀_{ĝ ∈ T} Nondet(ĝ);
4      result, sat_assignment ← Solve(φ);                          // run SAT solver
5      if result = UNSAT then
6          return SAFE, ν;
7      else
8          ϵ ← extract_CE(sat_assignment); // extract error trace from the sat.assignment
9          T ← T ∩ extract_calls(ϵ);          // filter out calls which do not affect SAT
10         if T = ∅ then
11             return BUG, ν, ϵ;
12         else
13             F̂_ν ← F̂_ν ∪ T;                          // unwind the calltree on demand
14             ν ← max_chain_length(F̂_ν);                          // update the depth
15 end
16 return TimeOut
```

calls from $\mathbb{T}$ are replaced by *true* (i.e., method `Nondet`). In total, ϕ encodes a program abstraction containing precise and over-approximated parts, conjoined by negation of an assertion π (line 3). Fig. 3a demonstrates a calltree of a program with a single recursive function called twice at the first iteration of the algorithm. In the example, $\hat{F}_\nu = \{\hat{f}_{main}, \hat{g}_1, \hat{h}_1, \hat{f}_1, \hat{f}_2\}$ (grey nodes) are encoded precisely, and $\mathbb{T} = \{\hat{f}_3, \hat{f}_2'\}$ (white nodes) are treated nondeterministically.

After the PBMC formula ϕ is constructed, the algorithm passes it to a SAT solver. If ϕ is satisfiable, and the SAT solver returns a satisfying assignment (line 7), function calls from $\mathbb{T}$ are considered as candidate calls to be refined. To refine, the satisfying assignment is used to restrict $\mathbb{T}$ on the calls, appeared along the error trace ϵ (i.e., in the satisfying assignment) (line 9). In the next iteration of the algorithm, the calls from $\mathbb{T}$ are encoded precisely in the updated PBMC formula. Technically, the algorithm extends $\hat{F}_\nu$ by adding function calls from $\mathbb{T}$ (line 13), as shown, for example, on Fig. 3b. There, $\hat{f}_2'$ appears along ϵ and therefore it has to be refined; $\hat{f}_3$ does not appear in ϵ, so it will be encoded nondeterministically. If $\mathbb{T} = \emptyset$ then no nondeterministically treated recursive calls were found along the error trace, so the real bug is found (line 11), and the algorithm terminates.

If the SAT solver proves unsatisfiability of ϕ then the program abstraction, and consequently the program itself, are safe (line 6). This case is represented on Fig. 3c. The final recursion depth ν is detected, and the algorithm terminates.

Theorem 1. *Given the program P and an assertion π, if Alg. 1 terminates with an answer* SAFE (BUG) *then π holds (does not hold) for P.*

Proof (Proof sketch). The proof is divided into two parts, for *SAFE* (line 6) and *BUG* (line 11) outputs of the algorithm (and respectively, the PBMC formula ϕ proven UNSAT or SAT).

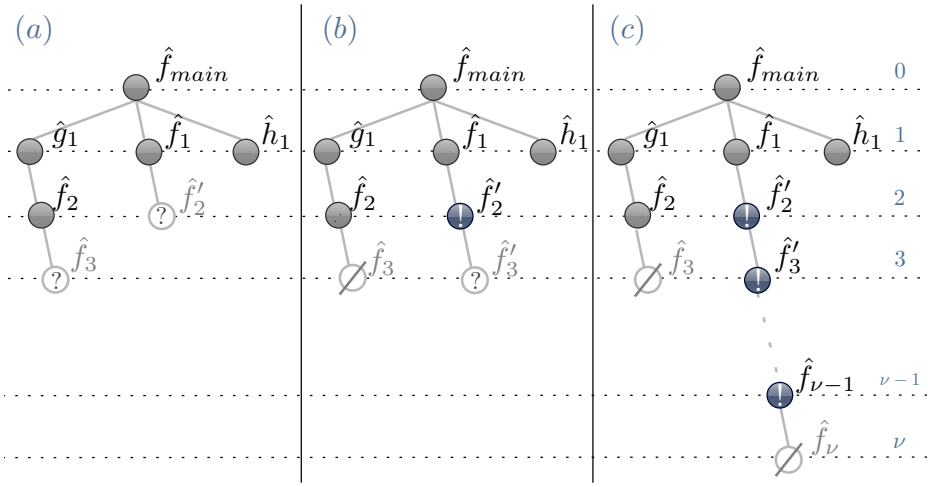

Figure 3: Illustration of the individual steps of the Alg. 1 on the example with a single recursive function f, called twice.
a) First iteration: $\hat{F}_\nu = \{\hat{f}_{main}, \hat{g}_1, \hat{h}_1, \hat{f}_1, \hat{f}_2\}$ (grey nodes) are encoded precisely, $\mathbb{T} = \{\hat{f}_3, \hat{f}'_2\}$ (white "?" nodes) are treated nondeterministically; the initial recursion depth is equal to 1.
b) Second iteration: solver returns SAT (corresponding to error trace $\epsilon = \{\hat{f}_{main}, \hat{f}_1, \hat{f}'_2\}$), set $\mathbb{T}$ is updated to contain only one function call ($\{\hat{f}'_2\}$ (black "!" nodes)). All calls from $\mathbb{T}$ are added to current $\hat{F}_\nu$. The current recursion depth is incremented, and equal to 2.
c) Final iteration: solver returns UNSAT or $\mathbb{T} = \emptyset$, the detected recursion depth is equal to $\nu - 1$.

Case ***SAFE***. In this case ϕ is unsatisfiable. The formula ϕ represents some abstraction of P which contains precise and over-approximated components (as described in section 3.1). Since every abstracted formula can be strengthened and turned into the corresponding precise encoding, and since unsatisfiability of a weaker formula implies unsatisfiability of a stronger formula, then the PBMC formula ϕ^{inline} encoding P without abstraction is also unsatisfiable, i.e., π holds.

Case ***BUG***. In this case, ϕ is satisfiable, and the satisfying assignment represents an error trace. At the same time, the algorithm did not detect any nondeterministically treated recursive function calls along the error trace (line 10). It means that π is indeed violated within the current recursion depth. □

Note on Termination. The algorithm is guaranteed to terminate within a given timeout when it finds an error or proves that the assertion holds. Similar to classical BMC, Alg. 1 terminates if the recursion depth is sufficient to disprove the assertion. Classical BMC can prove the assertion up to some fixed recursion depth, but the result might be incomplete if the recursion depth is insufficient. In contrast, by Theorem 1, if our algorithm does not yield a timeout, it guarantees that the detected recursion depth is complete to prove (disprove) the assertion. The other benefit of our algorithm is that it does not require the recursion depth to be given a priori, but instead it is detected automatically.

Based on our observations, termination of Alg. 1 depends on the termination of the recursive program it was applied to. For example, the program with one

Algorithm 2: Summarization in BMC with Automatic Detection of Recursion Depth

Input: Initial recursion depth ν; Program unwound ν times: $P_\nu = (\hat{F}_\nu, \hat{f}_{main}, child)$; Assertion to be checked: π; Set of summaries: $summaries$; $TimeOut$
Output: Verification result: $\{SAFE,\ BUG,\ TimeOut\}$; Error trace: ϵ
Data: ϕ: PBMC formula; $\mathbb{T}$: set of function calls; Ω: substitution scenario

1 $\phi \leftarrow \neg\pi$; // initialize ϕ
2 $\mathbb{T} \leftarrow \hat{F}_\nu \cup \{\hat{g} \notin \hat{F}_\nu \mid child(\hat{f}, \hat{g}), \hat{f} \in \hat{F}_\nu\}$; // unwind the calltree initially
3 $\Omega \leftarrow init$; // use (2) from Sect. 2.4 to create initial scenario
4 **while** $\neg$ $TimeOut$ **do**
5 $\phi \leftarrow \phi \wedge \bigwedge_{\hat{f}\in\mathbb{T}:\Omega(\hat{f})=inline}$ `CreateFormula`$(\hat{f}) \wedge \bigwedge_{\hat{g}\in\mathbb{T}:\Omega(\hat{g})=sum}$ `ApplySummaries`$(\hat{g}) \wedge \bigwedge_{\hat{h}\in\mathbb{T}:\Omega(\hat{h})=havoc}$ `Nondet`$(\hat{h})$; // add partitions to ϕ (*inline, summarize, havoc*)
6 $result,\ proof,\ sat\_assignment \leftarrow$`Solve`$(\phi)$;
7 **if** $result$ = UNSAT **then**
8 **foreach** $\hat{f} \in \mathbb{T}$: **do** // split $\phi \equiv \phi_{\hat{f}}^{subtree} \wedge \phi_{\hat{f}}^{env}$ as in Sect. 2.3
9 $summaries\ (\hat{f}) \leftarrow$ `Interpolate`$(proof, \hat{f})$;
10 **end**
11 **return** $SAFE$;
12 **else**
13 $\epsilon \leftarrow$ `extract_CE`$(sat\_assignment)$;
14 **if** $\{\hat{f} \in$ `extract_calls`$(\epsilon) \mid \Omega(\hat{f}) \neq inline\} = \emptyset$ **then**
15 **return** BUG, ϵ;
16 **else**
17 $\Omega \leftarrow$ `Refine`$(\Omega, \mathbb{T},$ `extract_calls`$(\epsilon))$; // use (3) from 2.4
18 $\mathbb{T} \leftarrow \mathbb{T} \cup \{\hat{g} \notin \mathbb{T} \mid child(\hat{f}, \hat{g}), \hat{f} \in \mathbb{T}, \Omega(\hat{f}) = inline\}$; // unwind the calltree on demand
19 **end**
20 **return** $TimeOut$

single recursive function from Fig. 1a terminates for any values of input data. The recursion termination condition, `¬(a < 10)` defines the upper bound 10 for the value of `a`, and at the same time the function `f` monotonically increments the value of `a`. Hence, the recursive function `f` is called a fixed number of times and the program eventually terminates. Clearly, for complete analysis of this program it is enough to consider the maximum possible number of recursive function calls for every initial value of `a` which in this example is equal to 5. At the same time, it introduces an upper bound for the size of the constructed PBMC formula which is a sufficient condition to the SAT solver to terminate while solving it.

3.2 Optimizations and Applications of Alg. 1

Incremental Formula Construction and Refinement. Possible optimizations of Alg. 1 are 1) the incremental construction of the PBMC formula ϕ and 2) more efficient handling of a set of the refinement candidates, $\mathbb{T}$.

In the first optimization, ϕ is created in an incremental manner. At each iteration, ϕ is not recomputed from scratch, but gets conjoined with new partitions. These partitions precisely encode the refined function calls from the set $\mathbb{T}$. In this manner the PBMC formula is updated at the beginning of each iteration.

In the second optimization, the set of refinement candidates $\mathbb{T}$ is merged with the whole set of unwound function calls $\hat{F}_\nu$. Instead of handling those two sets,

it is enough to handle one. To distinguish function calls which were present in $\mathbb{T}$ from the others present in $\hat{F}_\nu$ the substitution scenario Ω is used.

Summarization. The proposed algorithm for recursion depth detection can be exploited for efficient incremental program verification (i.e., verification of the same program with respect to different assertions [19].[5] In this setting, function summaries are computed by means of Craig Interpolation.

Alg. 2 shows how the optimized Alg. 1 can be integrated with summarization-based verification. Interpolating procedure (line 9), that employs the PBMC formula ϕ and its proof of unsatisfiability, is run after each assertion is proven. The use of summaries makes the verification more flexible. Instead of treating recursive function calls nondeterministically, the algorithm might apply existent summaries, thus making entire program abstraction more accurate. Moreover, the use of substitution scenario (line 5)[6] enables summarization of any (not necessarily recursive) function calls.

4 Experimental Evaluation

We implemented the automatic Recursion Depth Detection (RDD) and Summarization-based RDD (SRDD) inside of the BMC tool `FunFrog` [18] and make its binary (`FunFrog+(S)RDD`) available[7]. `FunFrog` supports interpolation-based function summarization for C programs and uses the SAT-solver `PeRIPLO` [17] for solving propositional formulas, proof reduction and interpolation. `FunFrog` follows CProver's[8] paradigm. In particular, it accepts a precompiled `goto-binary`, a representation of the C program in an intermediate `goto-cc` language, and runs the analysis on it.

We evaluated the new algorithms on a set of various recursive C programs (taken from the `SVCOMP'14`[9] set (`Ackermann X McCarthy`, `GCD`, `EvenOdd`), obtained from industry[10] (`P2P_Joints X`), crafted by USI students for evaluation of interpolation-based abstractions). We provide two verification scenarios to evaluate the algorithms. In the first one, `FunFrog+RDD` verifies a single assertion in each benchmark and detects the recursion depth. In the second one, `FunFrog+SRDD` incrementally verifies a set of assertions and reuses function summaries between its checks. In our experiments loop handling was done by means of CProver (see Sect. 5 for more details).

[5] Recall that the analysis in [19] is restricted to programs, unwound fixed number of times (i.e., without recursion).

[6] For simplicity, line 5 shows the construction of the whole PBMC formula ϕ. In practice, it does not recompute partitions of ϕ constructed in the previous iteration.

[7] `http://www.inf.usi.ch/phd/fedyukovich/funfrog_srdd.tar.gz`

[8] `http://www.cprover.org`

[9] `http://sv-comp.sosy-lab.org/2014/`

[10] in scope of FP7-ICT-2009-5 — project PINCETTE 257647

benchmark				FunFrog+RDD											FunFrog	CBMC
				In≡1					1 < In < ν			In≡ ν				
name	#R	T	Result	In	Time	#It	ν	#Calls	In	Time	#It	In	Time	#It	Time	Time
Array A	5	a	SAFE	1	664.02	15	15	75	10	513.986	6	15	121.381	1	3600+	3600+
Array B	12	a	SAFE	1	777.432	24	24	71	2	1781.92	23	24	3600+	—	3600+	3600+
Array C	3	a	SAFE	1	1113.68	27	16	106	14	991.724	3	16	557.281	1	3600+	3600+
Ackermann A	2	b	SAFE	1	55.758	34	20	2169	7	3493.64	10	20	3600+	—	3600+	3600+
Ackermann B	2	b	BUG	1	56.772	30	17	1942	7	3547.29	10	17	3600+	—	3600+	3600+
Alternate A	2	c	SAFE	1	35.068	50	50	100	30	22.206	20	50	0.902	1	3600+	3600+
Alternate B	2	c	BUG	1	92.314	77	77	154	50	53.315	28	77	1.681	1	3600+	3600+
Multiply	10	a	SAFE	1	710.517	110	10	110	7	569.559	4	10	226.659	1	3600+	3600+
InterleaveBitsRec	1	a	SAFE	1	150.053	33	33	33	15	125.241	19	33	8.188	1	3600+	3600+
BitShiftRec A	1	a	SAFE	1	128.074	64	64	64	20	13.416	45	64	2.413	1	3600+	3600+
BitShiftRec B	2	b	SAFE	1	65.537	12	12	4285	3	65.399	10	12	3600+	–	3600+	3600+
P2P_Joints A	1	a	SAFE	1	1234.71	4	4	4	2	1195.31	3	4	1092.26	1	3600+	3600+
P2P_Joints B	1	a	BUG	1	1266.38	4	4	4	2	1222.11	3	4	1120.03	1	3600+	3600+

Table 1: Verification statistics for various BMC tools with and without automated detection of recursion depth.

4.1 Evaluating RDD

Table 1 summarizes the verification statistics of a set of benchmarks with different types (**T**) of recursion (a - single recursion, b - multiple recursion, c - indirect recursion). The number of recursive functions present in each benchmark is depicted in the column marked #**R**. Each benchmark was verified using `CBMC`, `FunFrog`[11] without recursion depth detection and 3 different versions of `FunFrog+RDD`. The first configuration of `FunFrog+RDD` performs the algorithm with the initial recursion depth set to 1 (denoted as $\mathbf{In} \equiv \mathbf{1}$ in the table), detects recursion depth (ν) and also reports the number of unwound recursive calls as #**Calls**. Then, in purpose of comparison, the second and the third configurations perform the same algorithm with the another values of the initial recursion depths ($\mathbf{1 < In < \nu}$ and $\mathbf{In} \equiv \nu$ respectively). For each experiment, we report total verification time (in seconds) and a number of iterations of `FunFrog+RDD` (#**It**). The verification results (SAFE/BUG) were identical for experiments with all configurations and we placed them in the table in the section describing the benchmarks.

Notably, for all different types of recursion, the experiments with `CBMC` and pure `FunFrog` failed as they reached the timeout (**3600+**) of 1 hour without producing the result. This in general was not a problem for any of the experiments when `FunFrog+RDD` was used. We compare different configurations of `FunFrog+RDD` in order to demonstrate possible behaviors of `FunFrog+RDD` depending on the structure of benchmarks. The benchmarks `Multiply`, `Alternate A/B`, `Array A/C`, `InterleaveBitsRec` and `BitShiftRec A` witness the overhead of the procedure. In `InterleaveBitsRec` and `BitShiftRec A` there is a single recursive function called one time; in `Multiply` and `Alternate A/B` there are several recursive calls requiring the same recursion depth; in `Array A` and `Array C` there are several recursive calls requiring different, but relatively close recursion depths. That is, if we compare the first configuration with the third one, we can

[11] `CBMC` and `FunFrog` were run with default parameters

benchmark					FunFrog+RDD			FunFrog+SRDD				
name	#R	T	Result	ν	In	TotalTime	#It	In	#A	TotalTime	ItpTime	#It
Arithm	1	a	SAFE	100	1	128.47	100	1	20	9.676	2.036	119
McCarthy	2	b	SAFE	11	1	3600+	—	1	5	10.495	4.859	24
GCD	3	b	SAFE	11	1	145.381	64	1	4	54.185	0.409	37
EvenOdd	2	c	SAFE	25	1	38.621	50	1	8	27.99	4.49	82
P2P_Joints C	1	a	SAFE	4	1	1531.38	4	1	4	1151.72	68.10	4
P2P_Joints D	1	a	SAFE	4	1	1192.28	4	1	4	1089.04	87.08	4

Table 2: Verification statistics of FunFrog+RDD and FunFrog+SRDD

see that such overhead exists. The first configuration takes more time to complete verification than the second one, and the second configuration takes more time to complete verification than the third one. This is because FunFrog+RDD executes more iterations in the first configuration than in the second one and more iterations in the second configuration than in the third one. Again, the difference and the advantage is in the fact that the first and the second configurations do not know the recursion depth needed for verification and the third one gets it provided (as an initial recursion depth for FunFrog+RDD). Therefore, for the third configuration it is always enough to execute one iteration.

The benchmarks Array B, Ackerman A/B and BitShiftRec B show the opposite behavior, where the first configuration takes less time to complete than the second and the third ones. These cases demonstrate the benefits of using *minimality* feature of the FunFrog+RDD, since they require different recursion depths for each recursive function call appearing in the code. In all configurations we specify **In** by a fixed number which may fit well some of the recursive calls, but for other ones it may be bigger than needed. In this case, FunFrog+RDD creates unnecessary PBMC partitions, blows up the formula and consequently slows down the verification process. While using $\mathbf{In} = 1$, incremental unwinding automatically finds depths for each recursive function call. It means that for such cases the new approach for BMC not only detects the recursion depth sufficient for verification but that it also performs it efficiently and allows to slice out parts of the system which are redundant for verification purpose.

Interesting results are demonstrated by experimentation with the industrial benchmark P2P_Joints A/B. It contains expensive nonlinear computations, a complex calltree structure with relatively trivial recursion requiring unrolling 4 times. The experiments show that the difference in timings between different FunFrog+RDD configurations is minor.

4.2 Evaluating SRDD

Another set of experiments of verifying recursive programs by applying FunFrog+SRDD is summarized in Table 2. There are two configurations of FunFrog compared in the table. The first one, **FunFrog+RDD**, is similar to the first configuration in Table 1. The second one, **FunFrog+SRDD**, is SRDD driven by *assertion decomposition*.

We explain the idea of assertion decomposition on the example from Fig. 1. The assertion `assert(y >= 0)` ($A1$) can be used to derive a set the following *helper*-assertions `assert(x < 5 || y >= 0)` ($A2$), `assert(x < 7 || y >= 0)` ($A3$) and so on. It is clear that if $A1$ holds, then both $A2$ and $A3$ hold as well; and if $A2$ holds then $A3$ holds as well. We will say that $A3$ is *weaker* than $A2$, and $A2$ is *weaker* than $A1$.

In this experiment, we derive helper-assertions (number of them is denoted **#A** in the table) by guessing values of the input parameters of recursive functions, then order assertions by strength and begin verification from the weakest one. If the check succeeds, the summaries of all (even recursive) functions are extracted. They will be reused in verification of stronger assertions. This procedure is repeated until the original assertion is proven valid. We summarize total timings (**TotalTime**) for verification of each weaker assertion, which includes the timings for interpolation (**ItpTime**).

For all benchmarks in the table, `FunFrog+SRDD` outperforms `FunFrog+RDD`. Technically, it means that checking a single assertion may be slower than checking itself and also several other assertions.[12] The strongest result, we obtained, is verifying a well-known `McCarthy` function. Running `FunFrog+SRDD` for it takes around 10 seconds, while `FunFrog+RDD`, pure `FunFrog` and `CBMC` exceed timeout. Notably, the interpolation may take up to a half of whole verification time. In some cases, summarization increases the number of iterations. But in total, `FunFrog+SRDD` remains more efficient that `FunFrog+RDD`.

5 Related Work

To the best of our knowledge, there is very little support for computing recursion depths in BMC algorithms. One of the most successful BMC tools, `CBMC` [3], attempts to find unwinding recursion depths using constant propagation. This approach works only if the number of recursive calls is explicitly specified in the source code (i.e., as a constant number in a termination condition of a recursive call). If it cannot be detected by constant propagation, the tool gets into an infinite loop and fails to complete verification. `CBMC` also supports explicit definition of a recursion depth ν which may lead to incomplete verification results. In order to check correctness of the current unwinding, `CBMC` inserts and checks so called *unwinding assertions*. If all unwinding assertions hold, the currently used recursion depth is sufficient. If there is a violated unwinding assertion, the current recursion depth has to be increased. To our knowledge, `CBMC` does not have the refinement procedure and error trace analysis to make the recursion depth detection complete.

The idea of processing function calls on demand was also researched by [10] in the tool `Corral`. The method, called *stratified inlining*, relies on substituting bodies of function calls by summaries, and checking the resulting program using a theorem prover. If the given level of abstraction is not accurate enough, the

[12] A reader can find all these benchmarks with already inserted helper-assertions at `http://www.inf.usi.ch/phd/fedyukovich/funfrog_srdd.tar.gz`

algorithm refines function calls in a similar way to our refinement. Despite some similarity to Alg. 1, `Corral` relies on the external tool [7] to generate function summaries. In contrast, our method automatically generates summaries using Craig Interpolation inside Alg. 2 after an assertion is successfully checked, and use already constructed summaries to check other assertions.

There are techniques designed to deal with recursion. For instance, [20] is able to verify recursive programs in milliseconds, but it is limited only to functional programs. BMC, in contrast, is not designed to deal with recursion, but it has been applied to a wide range of verification tasks. `FunFrog+(S)RDD` itself is not a standalone recursive model checker, but an extension of the existent SAT-based BMC tool. In our previous work [18], it was already shown applicable to verify industrial-size programs, supporting complete ANSI C syntax. Conversion to SAT formulas allows to perform bit-precise checks, i.e., verify assertions in the programs using bitwise operators.

Craig Interpolation is applicable to verification of recursive programs in a rather different scenario. In `Whale` [1], it is used to guess summaries generated from under-approximations of the function bodies behavior. Unfortunately, the tool is not available for use, so we are unable to compare it with `FunFrog+(S)RDD`.

k-induction [5,15] is another under-approximation-driven technique for checking recursion. First, it proves an induction base (i.e., that there is no assertion violation in the unwinding chain with the length k). Then, if successful, it proves an induction step (i.e., whenever the assertion holds in an unwinding chain with the length k, it also holds in the unwinding chain with the length $(k+1)$). Finally, the approach is able to find an inductive invariant, which can be treated as function summary. To our knowledge, there is no incremental model checker based on k-induction which (re-)uses function summaries.

The overview of other summarization approaches to program analysis can be found in our earlier work published at [19].

6 Conclusion and Future Work

This paper presented the new approach to automatically detect recursion depths for BMC and applies it to function summarization-based approaches to model checking. In principle, a similar idea may be applied to solve the problem of loop bound detection where an algorithm abstracts away loop bodies and iteratively refines one more body at a time. One can develop such algorithm in future. We believe, there is a strong mapping between program termination and analysis termination which can be investigated in future. In cases of multiple recursion, the algorithm may be improved by using SAT solvers with support for Minimal SAT. The approach of the summarization-based BMC might be extended to support SMT theories. This way, the analysis in general might become more efficient, but will lose bit-precision.

Acknowledgments. We thank Antti Hyvärinen for his notable contribution during the work on this paper.

References

1. Albarghouthi, A., Gurfinkel, A., Chechik, M.: Whale: An interpolation-based algorithm for inter-procedural verification. In: VMCAI. LNCS, vol. 7148, pp. 39–55. Springer-Verlag (2012)
2. Biere, A., Cimatti, A., Clarke, E.M., Zhu, Y.: Symbolic Model Checking without BDDs. In: TACAS '99. LNCS, vol. 1579, pp. 193–207. Springer-Verlag (1999)
3. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: TACAS. pp. 168–176. LNCS, Springer-Verlag (2004)
4. Craig, W.: Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. In: J. of Symbolic Logic. pp. 269–285 (1957)
5. Donaldson, A.F., Haller, L., Kroening, D., Rümmer, P.: Software Verification Using k-Induction. In: SAS. pp. 351–368. LNCS, Springer-Verlag (2011)
6. Fedyukovich, G., Sery, O., Sharygina, N.: eVolCheck: Incremental Upgrade Checker for C. In: TACAS. LNCS, vol. 7795, pp. 292–307. Springer-Verlag (2013)
7. Flanagan, C., Leino, K.R.M.: Houdini, an Annotation Assistant for ESC/Java. In: FME. pp. 500–517. LNCS, Springer-Verlag (2001)
8. Graf, S., Saïdi, H.: Construction of Abstract State Graphs with PVS. In: Computer Aided Verification, CAV '97. pp. 72–83. LNCS, Springer-Verlag (1997)
9. Ivancic, F., Yang, Z., Ganai, M.K., Gupta, A., Ashar, P.: Efficient SAT-based bounded model checking for software verification. In: Theor. Comput. Sci. vol. 404, pp. 256–274 (2008)
10. Lal, A., Qadeer, S., Lahiri, S.K.: A solver for reachability modulo theories. In: CAV. LNCS, vol. 7358, pp. 427–443. Springer-Verlag (2012)
11. McMillan, K.L.: Applications of Craig Interpolation in Model Checking. In: TACAS. pp. 1–12. LNCS, Springer-Verlag (2005)
12. McMillan, K.L.: Lazy abstraction with interpolants. In: Computer Aided Verification (CAV '06). pp. 123–136. LNCS, Springer-Verlag (2006)
13. McMillan, K.L.: Lazy annotation for program testing and verification. In: Computer Aided Verification (CAV' 10). pp. 104–118. LNCS, Springer-Verlag (2010)
14. Merz, F., Falke, S., Sinz, C.: LLBMC: Bounded Model Checking of C and C++ Programs Using a Compiler IR. In: VSTTE. LNCS, vol. 7152, pp. 146–161. Springer-Verlag (2012)
15. Morse, J., Cordeiro, L., Nicole, D., Fischer, B.: Handling Unbounded Loops with ESBMC 1.20 - (Competition Contribution). In: TACAS. LNCS, vol. 7795, pp. 619–622. Springer-Verlag (2013)
16. Pudlák, P.: Lower bounds for resolution and cutting plane proofs and monotone computations. In: Journal of Symbolic Logic. vol. 62, pp. 981–998 (1997)
17. Rollini, S.F., Alt, L., Fedyukovich, G., Hyvärinen, A.E.J., Sharygina, N.: PeRIPLO: A Framework for Producing Effective Interpolant-based Software Verification. In: LPAR. LNCS, vol. 8312, pp. 683–693. Springer-Verlag (2013)
18. Sery, O., Fedyukovich, G., Sharygina, N.: FunFrog: Bounded Model Checking with Interpolation-based Function Summarization. In: ATVA. LNCS, vol. 7561, pp. 203–207. Springer-Verlag (2012)
19. Sery, O., Fedyukovich, G., Sharygina, N.: Interpolation-based Function Summaries in Bounded Model Checking. In: HVC. LNCS, vol. 7261, pp. 160–175. Springer-Verlag (2012)
20. Unno, H., Terauchi, T., Kobayashi, N.: Automating relatively complete verification of higher-order functional programs. In: POPL. pp. 75–86. ACM (2013)

A Types of recursion

Fig. 4 demonstrates different types of possible recursive function calls up to the depth ν.

Fig. 4a shows an example with a single recursive function f called two times, once from function g and once from function f_{main}. In this example, the calltree contains two chains of calls of function f: the first one consisting of one function call $\{\hat{f}_2\}$, the other consisting of ν calls: $\{\hat{f}_1, \hat{f}'_2, ..\hat{f}_\nu\}$, where the numbers 1 and ν are recursion depths.

Fig. 4b shows an example with a recursive function f called multiple times from itself (in the example, it is called 2 times). There are many chains of function calls possible for such scenario, and every one consists of at most ν calls of f, as demonstrated by a sample unwinding in the example. Notably, their unwinding depths can be different (and our algorithm will be able to detect the longest ones and stop exploring the chains for which the smaller depth is sufficient for verification).

Fig. 4c shows an example with indirect recursive functions f and g, such that each function is called not by itself, but by another function that it called. In the example, both f an g are unwound *at most* $\lfloor \frac{\nu}{2} \rfloor$ times (i.e., ν times altogether).

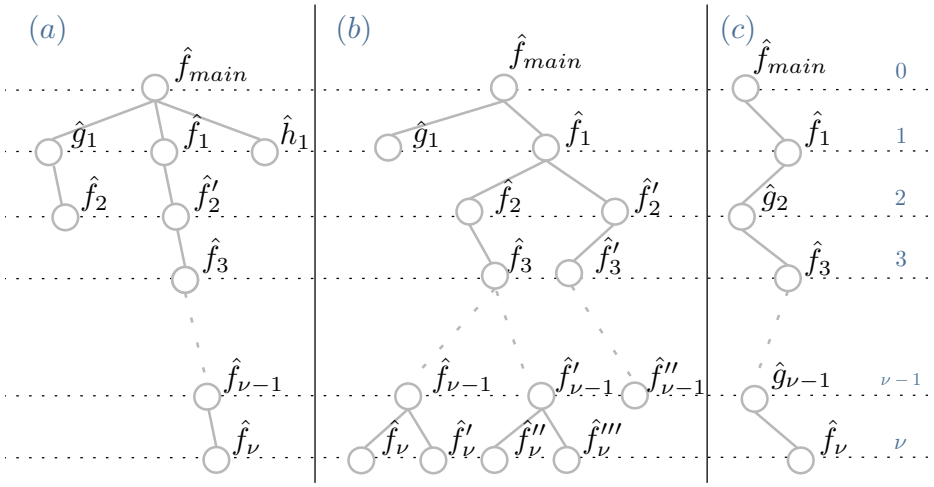

Figure 4: A program calltree with recursive functions unwound at most ν times: a) single recursion; b) multiple recursion; c) indirect recursion

Completeness and decidability results for hybrid(ised) logics

Renato Neves[1], Manuel A. Martins[2], and Luís S. Barbosa[1]

[1] HASLab INESC TEC & Univ. Minho
{nevrenato, lsb}@di.uminho.pt
[2] CIDMA - Dep. Mathematics, Univ. Aveiro
martins@ua.pt

Abstract. Adding to the modal description of transition structures the ability to refer to specific states, hybrid(ised) logics provide an interesting framework for the specification of reconfigurable systems. The qualifier 'hybrid(ised)' refers to a generic method of developing, on top of whatever specification logic is used to model software configurations, the elements of an hybrid language, including nominals and modalities. In such a context, this paper shows how a calculus for a hybrid(ised) logic can be generated from a calculus of the base logic and that, moreover, it preserves soundness and completeness. A second contribution establishes that hybridising a decidable logic also gives rise to a decidable hybrid(ised) one. These results pave the way to the development of dedicated proof tools for such logics used in the design of reconfigurable systems.

1 Introduction

1.1 Motivation

The need to master ubiquitous and increasingly complex software systems, often of a safety–critical nature, has brought proof and verification to a central place in Computer Science and Software Engineering. Logics, as formal reasoning frameworks, provide tools for a rigorous specification (and analysis) of software systems, as opposed to more conventional practices in software development which are often pre-scientific and unable to prove the absence of error designs.

Ideally, the working software engineer seeks for logics that can effectively provide "yes–or–no" answers to queries regarding properties of the system (*i.e.* *decidable* logics), as well as logics with a calculus providing enough syntactic rules to derive falsehood from any false statement (*i.e.* a *complete* calculus). The engineer also looks for logics with the right expressive power to specify the system at hand, a job made difficult by the complex and heterogeneous nature of current software systems which typically require a number of different logics to be suitably specified. For example, some form of equational logic may be used for data type specifications, while transitional behaviour my resort to a modal or temporal logic and fuzzy requirements may become in order to express contextual constraints. Actually, this justifies the quest for methodologies in which a

specification framework can be tailored by combining whichever logics are found suitable to deal with the different nature of the requirements in presence. As Goguen and Meseguer put it in a landmark paper [11],

> *"The right way to combine various programming paradigms is to discover their underlying logics, combine them, and then base a language upon the combined logic."*

This line of research has been particularly active for the last twenty years. Finger and Gabbay, for example, showed in [9] how to add a temporal dimension to an arbitrary logic, and proved that decidability and completeness is preserved along this process. Baltazar [2] did similar work but with respect to adding a probabilistic dimension. Other, similar results include *e.g.* [6], [7], as well as a *hybridisation* method [14], in whose development the current authors have been involved, and constitutes the starting point of the work reported in the sequel.

1.2 Context

Essentially hybridisation turns a given logic, defined as an *institution*, into a hybrid logic, a brand of modal logics that adds to the modal description of transition structures the ability to refer to specific states (*cf.* [1, 3]). This paves the way to an expressive framework, proposed in [13], for the specification of *reconfigurable* systems, *i.e.*, systems which may evolve through different execution modes, or configurations, along their lifetime. Specification proceeds in two steps:

- *globally* the system's dynamics is represented by a transition structure described in a hybrid language, whose states correspond to possible configurations;
- *locally* each state is endowed with a structure modelling the specification of the associated configuration.

The logic used locally, *i.e.* the one to be hybridised, depends on the application requirements. Typical candidates are equational, partial algebra or first-order logic (FOL), but one may equally resort to multivalued logics or even to hybrid logic itself equipping, in the last case, each state with another (local) transition system. Verification resorts to a parametrised translation to FOL (developed in [14] and [15]), but at the cost of losing decidability and adding extra complexity.

The generic character of this hybridisation process is achieved through its rendering in the context of institution theory [10]. Such a theory formalises the essence of what a logical system actually is, by encompassing syntax, semantics and satisfaction. However, its classical definition, the one in which the hybridisation method is based, does not include an abstract structure to represent a logic calculus. The problem was addressed in [8] with the introduction of π–*institutions*, and, more recently, in [5] with the notion of an *institution with proofs*, a more general version of the previous work.

1.3 Contributions and roadmap

This paper starts by recasting the hybridisation method in the theory of institutions with proofs, which makes possible the systematic generation of a calculus when hybridising a given logic.

Then, we prove that, under certain conditions, this method preserves decidability, and furthermore that the generated calculus is sound and complete whenever the one corresponding to the base logic is. Those are the paper's main contributions. Besides their theoretical relevance, from a pragmatic point of view they pave the way to the development of effective verification algorithms.

The paper is organised as follows. Institutions with proofs are briefly reviewed in Section 2. Then, Section 3 introduces the generation of an hybrid calculus from a base one. Section 4 establishes decidability and completeness. Finally, Section 5 concludes the paper and hints at future lines of research.

2 Background

We first recall the notion of an institution [10]. As already mentioned, it formalises the essence of a logical system, encompassing syntax, semantics and satisfaction. Put forward by J. Goguen and R. Burstall in the late seventies, its original aim was to develop as much as Computer Science as possible in a general uniform way independently of particular logical systems. This has now been achieved to an extent even greater than originally thought, with the theory of institutions becoming the most fundamental mathematical theory underlying algebraic specification methods, and also increasingly used in other areas of Computer Science. Formally,

Definition 1. *An institution is a tuple* $(Sign^I, Sen^I, Mod^I, (\models^I_\Sigma)_{\Sigma \in |Sign^I|})$, *where:*

- $Sign^I$ *is a category whose objects are signatures and arrows signature morphisms,*
- $Sen^I : Sign^I \to \mathbb{S}et$, *is a functor that, for each signature* $\Sigma \in |Sign^I|$, *returns a set of sentences over* Σ,
- $Mod^I : (Sign^I)^{op} \to \mathbb{C}at$, *is a functor that, for each signature* $\Sigma \in |Sign^I|$, *returns a category whose objects are models over* Σ,
- $\models^I_\Sigma \subseteq |Mod^I(\Sigma)| \times Sen^I(\Sigma)$, *or simply* $\models$, *if the context is clear, is a satisfaction relation such that, for each signature morphism* $\varphi : \Sigma \to \Sigma'$,

$$Mod^I(\varphi)(M') \models^I_\Sigma \rho \text{ iff } M' \models^I_{\Sigma'} Sen^I(\varphi)(\rho), \text{ for any}$$

$M' \in |Mod^I(\Sigma')|$ *and* $\rho \in Sen^I(\Sigma)$. *Graphically,*

$$\begin{array}{ccc} \Sigma & Mod^I(\Sigma) & \xrightarrow{\models^I_\Sigma} Sen^I(\Sigma) \\ {\scriptstyle\varphi}\downarrow & {\scriptstyle Mod^I(\varphi)}\uparrow & \downarrow{\scriptstyle Sen^I(\varphi)} \\ \Sigma' & Mod^I(\Sigma') & \xrightarrow[\models^I_{\Sigma'}]{} Sen^I(\Sigma') \end{array}$$

Intuitively, this property means that satisfaction is preserved under change of notation.

Definition 2. *Consider an institution I and signature $\Sigma \in |Sign^I|$. We say that a sentence $\rho \in Sen^I(\Sigma)$ is Σ–valid (or simply, valid) if for each model $M \in |Mod^I(\Sigma)|$, $M \models^I_\Sigma \rho$. Usually we prefix such sentences by $\models^I_\Sigma$ or, simply by $\models^I$ or just $\models$.*

Definition 3. *An institution I has the* negation property *if, for any signature $\Sigma \in |Sign^I|$ and sentence $\rho \in Sen^I(\Sigma)$, there is a sentence, $\neg\rho \in Sen^I(\Sigma)$, such that for any model $M \in |Mod^I(\Sigma)|$, $M \models^I_\Sigma \rho$ iff $M \not\models^I_\Sigma \neg\rho$.*

If this property holds, satisfiability of sentences may be rephrased as follows,

Definition 4. *Consider institution I with the* negation property *and a signature $\Sigma \in |Sign^I|$. For any sentence $\rho \in Sen^I(\Sigma)$,*

$$\rho \textit{ is } \Sigma\text{–unsatisfiable iff } \neg\rho \textit{ is } \Sigma\text{–valid.}$$

Similarly,

Definition 5. *An institution I has the* explicit satisfaction property*, if for any signature $\Sigma \in |Sign^I|$ and sentence $\rho \in Sen^I(\Sigma)$, satisfiability of ρ entails the existence of a model $M \in |Mod^I(\Sigma)|$ such that $M \models^I_\Sigma \rho$.*

Note that this last property holds in the most common logics used in specification, *e.g.*, propositional, fuzzy, equational, partial and first-order.

Definition 6. *An institution I has the* conjunction property *if, for any signature $\Sigma \in |Sign^I|$ and sentences $\rho, \rho' \in Sen^I(\Sigma)$, there is sentence $\rho \wedge \rho' \in Sen^I(\Sigma)$, such that for any model $M \in |Mod^I(\Sigma)|$, $M \models^I_\Sigma \rho \wedge \rho'$ iff $M \models^I_\Sigma \rho$ and $M \models^I_\Sigma \rho'$*

Note that with the conjunction property we are able to define a sentence $(\rho \wedge \neg\rho) \in Sen^I(\Sigma)$, denoted by $\bot$, that is not satisfied by any model of $|Mod^I(\Sigma)|$.

An institution for which both the negation and conjunction properties hold, is said to have the typical boolean connectives.

In order to better grasp this rather abstract concept of an institution let us analyse some typical examples.

Example 1. Many sorted first order logic (FOL)

- SIGNATURES. $Sign^{\mathrm{FOL}}$ is a category whose objects are triples (S, F, P), consisting of a set of sort symbols S, a family, $F = (F_{w\to s})_{w\in S^*, s\in S}$, of function symbols indexed by their arity, and a family, $P = (P_w)_{w\in S^*}$, of relational symbols also indexed by their arity.
  A signature morphism in this category is a triple $(\varphi_{st}, \varphi_{op}, \varphi_{rl}) : (S, F, P) \to (S', F', P')$ such that if $\sigma \in F_{w\to s}$, then $\varphi_{op}(\sigma) \in F'_{\varphi_{st}(w)\to\varphi_{st}(s)}$, and if $\pi \in P_w$ then $\varphi_{rl}(\pi) \in P'_{\varphi_{st}(w)}$.

– SENTENCES. For each signature object $(S, F, P) \in |Sign^{\mathrm{FOL}}|$, $Sen^{\mathrm{FOL}}\big((S, F, P)\big)$ is the smallest set generated by:

$$\rho \ni \neg\rho \mid \rho \wedge \rho \mid t = t \mid \pi(X) \mid \forall x : s \,.\, \rho'$$

where t is a term of sorts with the syntactic structure $\sigma(X)$ for $\sigma \in F_{w\to s}$ and X a list of terms compatible with the arity of σ. $\pi \in P_w$ and X is a list of terms compatible with the arity of π. Finally, $\rho' \in Sen^{FOL}\big((S, F \uplus \{x\}_{\to s}, P)\big)$. $Sen^I(\varphi)$, for φ a signature morphism, is a function that, given a sentence $\rho \in Sen^I\big((S, F, P)\big)$, replaces the signature symbols in ρ under the mapping corresponding to φ.

– MODELS. For each signature $(S, F, P) \in |Sign^{\mathrm{FOL}}|, Mod^{\mathrm{FOL}}\big((S, F, P)\big)$ is category with only identity arrows and whose objects are models with a carrier set $|M_s|$, for each $s \in S$; a function $M_\sigma : |M_w| \to |M_s|$, for each $\sigma_{w\to s} \in F_{w\to s}$; a relation $M_\pi \subseteq |M_w|$, for each $\pi \in P_w$.

– SATISFACTION. Satisfaction of sentences by models is the usual Tarskian satisfaction.

Example 2. Equational logic (EQ)

The institution EQ is the sub-institution of FOL in which sentences are restricted to those of the type $\forall x : s \,.\, t = t'$

Example 3. Propositional logic (PL)

Institution PL is the sub–institution of FOL in which signatures with no empty set of sorts are discarded.

As seen above, no notion of a proof system is considered in the definition of an institution. This is a limitation if one is interested in logical systems with calculi, as is the case in this paper which aims at introducing the systematic generation of calculi for hybridised logics. To overcome this we resort to the following extended definition of an institution with proofs [5].

Definition 7. *An institution with proofs adds to the original definition a functor $Prf^I : Sign^I \to \mathbb{C}at$ such that, for each $\Sigma \in |Sign^I|$, $Prf(\Sigma)$ (called the category of Σ–proofs) has subsets of $Sen^I(\Sigma)$* (i.e., $|Prf(\Sigma)| = |\mathcal{P}\big(Sen^I(\Sigma)\big)|$) *as objects, and the corresponding proofs as arrows. The latter are preserved along signature morphisms. In addition, for $A, B \in |Prf^I(\Sigma)|$, if $A \subseteq B$ then there is an arrow $B \longrightarrow A$; if $A \cap B = \emptyset$ and there is $\Gamma \in |Prf^I(\Sigma)|$ such that $p : \Gamma \longrightarrow A$ and $q : \Gamma \longrightarrow B$, then there is a unique proof $\langle p, q\rangle$ making the following diagram to commute*

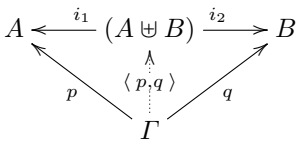

Note that the restrictions imposed to the proof arrows oblige Prf^I to follow the basic properties of a proof system. In particular, we have

1. *Reflexivity* (if $A \in \Gamma$, then $\Gamma \vdash A$) follows from the fact that $\{A\} \subseteq \Gamma$ and therefore $\Gamma \longrightarrow A$.
2. *Monotonicity* (if $\Gamma \vdash A$ and $\Gamma \subseteq \Delta$ then $\Delta \vdash A$), follows from composition of proofs, where $\Delta \longrightarrow \Gamma$ is given by inclusion and $\Gamma \longrightarrow A$ by the assumption.
3. *Transitivity* (if $\Gamma \vdash A$ and $\{\Delta, A\} \vdash B$ then $\Gamma \cup \Delta \vdash B$), follows from the product of disjoint sets, reflexivity and monotonicity,

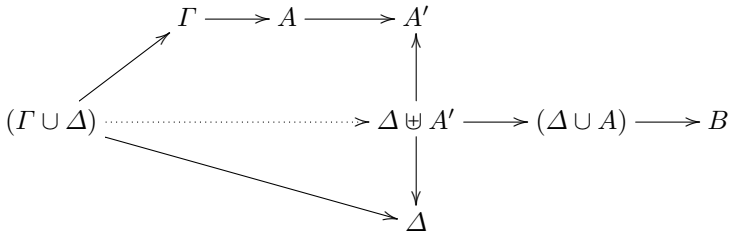

where $A' = A - (A \cap \Delta)$ ($A' \subseteq A$ and $(\Delta \cup A) \subseteq (\Delta \cup A')$).

Note that functor Prf^I distinguishes different proofs between the same pair of objects, as opposed to entailment systems[3]. In this work, however, we restrict ourselves to entailment systems in which $Prf^I(\Sigma)$ has at most one arrow for each pair of objects, *i.e.* that $Prf^I(\Sigma)$ is *thin*. Such restriction makes showing the uniqueness of $\langle p, q \rangle$ trivial. Also for the sake of simplicity, when a singleton set of sentences is presented in a proof arrow, we may drop the curly brackets.

Definition 8. *Let I be an institution with proof system Prf^I. We say that Prf^I is* sound *if, for any signature $\Sigma \in |Sign^I|$ and sentence $\rho \in Sen^I(\Sigma)$,*

$$\textit{if arrow, } \emptyset \longrightarrow \rho, \textit{ is in } Prf^I(\Sigma) \textit{ then } \models^I \rho.$$

Definition 9. *Let I be an institution with proof system Prf^I. We say that Prf^I is* complete *if, for any signature $\Sigma \in |Sign^I|$ and sentence $\rho \in Sen^I(\Sigma)$,*

$$\textit{if } \models^I \rho \textit{ then arrow, } \emptyset \longrightarrow \rho, \textit{ is in } Prf^I(\Sigma)$$

Hence, soundness and completeness of Prf^I entails the equivalence, for any signature $\Sigma \in |Sign^I|$ and sentence $\rho \in Sen^I(\Sigma)$,

$$\models^I \rho \quad \text{iff} \quad \emptyset \longrightarrow \rho \text{ is in } Prf^I(\Sigma)$$

We can now show that

Theorem 1. *If an institution I has classical boolean connectives, and a sound and complete calculus Prf^I, with the* reductio ad absurdum *property, then, for any signature, $\Sigma \in |Sign^I|$, and sentence, $\rho \in Sen^I(\Sigma)$,*

$$\rho \textit{ is satisfiable iff } \rho \longrightarrow \bot \textit{ is not in } Prf^I(\Sigma)$$

[3] Typically, in an entailment system $\Gamma \vdash A$ means that Γ derives (or entails) A.

Proof.

$$
\begin{array}{ll}
 & \models^I \rho \ \text{ iff } \ \emptyset \longrightarrow \rho \text{ is in } Prf^I(\Sigma) \\
\Leftrightarrow & \quad \{ \text{ defn. of satisfiability } \} \\
 & \neg\rho \text{ is unsat } \text{ iff } \ \emptyset \longrightarrow \rho \text{ is in } Prf^I(\Sigma) \\
\Leftrightarrow & \quad \{ \text{ soundness, completeness of } Prf^I(\Sigma) \text{ and } r.a.a \} \\
 & \neg\rho \text{ is unsat } \text{ iff } \ \neg\rho \longrightarrow \bot \text{ is in } Prf^I(\Sigma) \\
\Leftrightarrow & \quad \{ \text{ defn. of negation } \} \\
 & \rho \text{ is unsat } \text{ iff } \ \rho \longrightarrow \bot \text{ is in } Prf^I(\Sigma) \\
\Leftrightarrow & \quad \{ \text{ de Morgan's law} \} \\
 & \rho \text{ is sat } \text{ iff } \ \rho \longrightarrow \bot \text{ is not in } Prf^I(\Sigma)
\end{array}
$$

Corollary 1. *In the context of theorem 1, if I has the explicit satisfaction property, then*

$$
\begin{array}{ll}
 & \rho \textit{ is sat } \textit{ iff } \ \rho \longrightarrow \bot \textit{ is not in } Prf^I(\Sigma) \\
\Leftrightarrow & \quad \{ \textit{ explicit satisfaction property } \} \\
 & \rho \textit{ has a model } \textit{ iff } \ \rho \longrightarrow \bot \textit{ is not in } Prf^I(\Sigma)
\end{array}
$$

This last result will be essential in the sequel for proving completeness of hybridised logics.

3 Hybridisation of logics and their calculi

As mentioned before, the existence of software products that are built and maintained with respect to requirements of different nature calls for techniques that favour combination of logics. Hybridisation [14] was born in this context. It aims at providing a framework to specify reconfigurable systems, whose execution modes are described by whatever logic the engineer finds suitable, whereas the transition structure is expressed in a hybrid language.

From a point of view of verification, however, the engineer is not only interested in having a hybridised logic, but also, in a very pragmatic way, in its calculus. This section addresses such issue. It starts by revisiting hybridisation and then, through the notion of institutions with proofs, it shows how to lift the calculus in the base logic to its hybridised counterpart.

3.1 Hybridisation revisited

Definition 10. *The category $Sign^{\mathcal{H}}$ is the category $\mathbb{S}et \times \mathbb{S}et$ whose objects are pairs $(\ Nom, \Lambda\ )$, with Nom denoting a set of nominal symbols and, Λ, a set of modality symbols.*

Definition 11. *Given an institution* $I = (Sign^I, Sen^I, Mod^I, \models^I)$ *its hybridised version* $\mathcal{H}I = (Sign^{\mathcal{H}I}, Sen^{\mathcal{H}I}, Mod^{\mathcal{H}I}, \models^{\mathcal{H}I})$ *is defined as follows,*

- $Sign^{\mathcal{H}\mathcal{I}} = Sign^{\mathcal{H}} \times Sign^{\mathcal{I}}$,
- *given a signature* $(\Delta, \Sigma) \in |Sign^{\mathcal{H}\mathcal{I}}|$, $Sen^{\mathcal{H}\mathcal{I}}\big((\Delta, \Sigma)\big)$ *is the least set generated by*

$$\rho \ni \neg\rho \mid \rho \barwedge \rho \mid i \mid @_i\rho \mid \langle\lambda\rangle\rho \mid \forall x\ \rho' \mid \psi \mid \mathrm{A}\ \rho$$

  for i *a nominal,* λ *a modality,* $\psi \in Sen^{\mathcal{I}}(\Sigma)$ *and* $\rho' \in Sen^{\mathcal{H}I}\big((\Delta \uplus \{x\}, \Sigma)\big)$ *where* x *is a nominal. We use non standard boolean connectives* $(\neg, \barwedge)$[4] *in order to distinguish them from the boolean connectives that the base logic may have.*
- *given a signature* $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$, *a model* $M \in |Mod^{\mathcal{H}I}\big((\Delta, \Sigma)\big)|$ *is a triple* (W, R, m) *such that,*
  - W *is a non–empty set of worlds,*
  - R *is a family of relational symbols indexed by the modality symbols, such that for each* $\lambda \in \Lambda$ *(* where $\Delta = (-, \Lambda)$ *),* $R_\lambda \subseteq W \times W$,
  - *and* $m : W \to |Mod^I(\Sigma)|$,

  and for each $i \in Nom$, $(W, R, m)_i$ *is interpreted as a world in* W.
- *given a signature* $(\Delta, \Sigma) \in |Sign^I|$, *a model* $M = (W, R, m) \in |Mod^I\big((\Delta, \Sigma)\big)|$ *and a sentence* $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$, *the satisfaction relation is defined as,*

$$M \models^{\mathcal{H}I}_{(\Delta,\Sigma)} \rho \text{ iff } M \models^w \rho, \text{ for all } w \in W$$

  where,
  $M \models^w \neg\rho$ *iff* $M \not\models^w \rho$
  $M \models^w \rho \barwedge \rho'$ *iff* $M \models^w \rho$ *and* $M \models^w \rho'$
  $M \models^w i$ *iff* $M_i = w$
  $M \models^w @_i\rho$ *iff* $M \models^{M_i} \rho$
  $M \models^w \langle\lambda\rangle\rho$ *iff there is some* $w' \in W$ *such that* $(w, w') \in R_\lambda$ *and* $M \models^{w'} \rho$
  $M \models^w \mathrm{A}\rho$ *iff* $M \models^w \forall x\ @_x\rho$
  $M \models^w \psi$ *iff* $m(w) \models^I_\Sigma \psi$
  $M \models^w \forall x\ \rho$ *iff for all* M', $M' \models \rho$

for $(W, R, m) = M' \in |Mod^{\mathcal{H}I}\big((\Delta \uplus \{x\}, \Sigma)\big)|$ *a model expansion of* M, *with the only difference between them being the interpretation of nominal* x*: while it is defined in* M', *in* M *it is not.*

Note that sentence ρ being satisfiable means that there is a model $(W, R, m) = M \in |Mod^{\mathcal{H}I}\big((\Delta, \Sigma)\big)|$ such that $M \models^w \rho$ for some $w \in W$. Hence, hybridised logics do not have the explicit satisfaction property. One can, however, redefine the satisfaction relation in the hybridisation method to,

$$M \models^{\mathcal{H}I}_{(\Delta,\Sigma)} \rho \text{ iff } M \models^w \rho, \text{ for some } w \in W$$

which then provides to logics hybridised in this alternative way the explicit satisfaction property.

A weak hybridisation of an institution I, denoted by $\mathcal{H}'I$, is obtained as $\mathcal{H}I$, but the omission of syntax constructor $\forall x\ \rho$. The following decidability results are formulated with respect to weak hybridisation.

[4] implication ($\Rightarrow$) and biimplication ($\Leftrightarrow$) are built in the usual way.

3.2 Hybridising a calculus

We now present the hybridisation of calculi in the context of institutions with proofs. Let us assume that I has a proof system, *i.e.*, that Prf^I is well defined, and that, in particular, it is an entailment system, *i.e.*, Prf^I only defines thin categories. Then we define $Prf^{\mathcal{H}I}$ as follows:

For any $\big((Nom, \Lambda), \Sigma\big) \in |Sign^{\mathcal{H}I}|$,

1. for any $\rho \in Sen^I(\Sigma)$, if $\emptyset \longrightarrow \rho$ is in $Prf^I(\Sigma)$ then $\emptyset \longrightarrow \rho$ is in $Prf^{\mathcal{H}I}((Nom, \Lambda), \Sigma)$,
2. for any nominal $i, j \in Nom$, modality $\lambda \in \Lambda$, $\rho, \rho' \in Sen^{\mathcal{H}I}\big((Nom, \Lambda), \Sigma\big)$, proof arrows in Table 1 are in $Prf^{\mathcal{H}I}((Nom, \Lambda), \Sigma)$
3. finally, $Prf^{\mathcal{H}I}((Nom, \Lambda), \Sigma)$ has all the inclusion proof arrows and for each $A, B, \Gamma \in |Prf^{\mathcal{H}I}((Nom, \Lambda), \Sigma)|$ if $\Gamma \longrightarrow A$, $\Gamma \longrightarrow B$ then $\Gamma \longrightarrow A \cup B$.

$Prf^{\mathcal{H}I}$ is maintained thin in its construction process in order to have it as an entailment system.

Axioms

(CT)	All substitution instances of classical tautologies
$(Dist)$	$\emptyset \longrightarrow @_i(\rho \Rightarrow \rho') \Leftrightarrow (@_i\rho \Rightarrow @_i\rho')$
$(\bot)$	$\emptyset \longrightarrow @_i\bot \Rightarrow \bot$
(Scope)	$\emptyset \longrightarrow @_i@_j\rho \Rightarrow @_j\rho$
(Ref)	$\emptyset \longrightarrow @_i i$
(Intro)	$\emptyset \longrightarrow (i \mathbin{\&} \rho) \Rightarrow @_i\rho$
$(\Box E)$	$\emptyset \longrightarrow ([\lambda]\rho \mathbin{\&} \langle\lambda\rangle i) \Rightarrow @_i\rho$
$(\forall E)$	$\emptyset \longrightarrow \forall x\ \rho \Rightarrow \rho[i/x]$

Rules

(MP)	if $\emptyset \longrightarrow \rho$ and $\rho \longrightarrow \rho'$ then $\emptyset \longrightarrow \rho'$
$(N_@)$	if $\emptyset \longrightarrow \rho$ then $\emptyset \longrightarrow @_i\rho$
$(Name)$	if i does not occur free in ρ and $\emptyset \longrightarrow @_i\rho$ then $\emptyset \longrightarrow \rho$
$(\Box I)$	if i does not occur free in ρ, ρ' and $\emptyset \longrightarrow (\rho \wedge \langle\lambda\rangle i) \Rightarrow @_i\rho'$ then $\emptyset \longrightarrow \rho \Rightarrow [\lambda]\rho'$
$(\forall I)$	if i does not occur free in $\forall x\ \rho', \rho$ and $\emptyset \longrightarrow \rho \Rightarrow\ \rho'[i/x]$ $\emptyset \longrightarrow \rho \Rightarrow \forall x\ \rho'$

Table 1. Axioms and rules for $Prf^{\mathcal{H}I}$ from [3]

4 Decidability and completeness of hybridised logics

Decidability and completeness are properties that one usually looks for when researching a new logic. From a Computer Science perspective, they are essential as a basis for tool-supported proofs. Formally,

Definition 12. *Decidability of an institution I means that, for each signature $\Sigma \in |Sign^I|$ and sentence $\rho \in Sen^I(\Sigma)$, there is an effective algorithm able to tell if ρ is valid.*

After some preliminary work, we address first this definition in the context of hybridised logics.

4.1 Preliminaries

Recall that in the sequel we assume that the base institution I has the classical boolean connectives and the explicit satisfaction property. Furthermore, its calculus, Prf^I, is sound, complete and has the *reductio ad absurdum* property.

Notation 1. *Consider* $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$ *and* $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$. *Let* $B_\rho = \{\psi_1, \ldots, \psi_n\}$ *to denote the set of all maximal sentences,* $\psi_i \in Sen^I(\Sigma)$, *occurring in* ρ. *Then, the set of base sentences,* Ω_ρ, *denotes the least set such that for each* $a \in 2^{B_\rho}$,

$$(\chi_1 \wedge \cdots \wedge \chi_n) \in \Omega_\rho \subseteq Sen^I(\Sigma)$$

where

$$\chi_i = \begin{cases} \psi_i & \text{if } \psi_i \in a \\ \neg\psi_i & \text{if } \psi_i \notin a \end{cases}$$

Whenever suitable we abbreviate $(\chi_1 \wedge \cdots \wedge \chi_n)$ *to* χ, *and refer to components of* χ *as* χ_i.

Lemma 1. *For any model* $M \in |Mod^I(\Sigma)|$, M *satisfies exactly one of the sentences in* Ω_ρ.

Proof. Suppose that M fails to satisfy a sentence $\chi \in \Omega_\rho$. This only happens when at least one member of χ is not satisfied by M. By the definition of Ω_ρ we know that Ω_ρ has another sentence χ' which negates all the failed components in χ and therefore M must satisfy χ'.

Suppose that M satisfies a sentence $\chi \in \Omega_\rho$. Clearly, by the definition of Ω_ρ any other sentence $\chi' \in \Omega_\rho$ must negate at least one of the components of χ. Since M cannot satisfy a component and its negation, χ' cannot be satisfied by M.

Notation 2. *If* Ω_ρ *is not empty, Lemma 1 allows the use of notation* Ω_ρ^M *to denote the sentence in* Ω_ρ *which is satisfied by a model* $M \in |Mod^I(\Sigma)|$.

Next, in order to take advantage of the well known decidability and completeness results for hybrid propositional logic, $\mathcal{H}PL$, we define a function between $\mathcal{H}I$ and $\mathcal{H}PL$ sentences,

Definition 13. *Consider a signature* $(\Delta, \Sigma) \in |Sen^{\mathcal{H}I}|$, *a sentence* $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$, *and a PL signature Prop that, for each* $\psi_i \in Sen^I(\Sigma)$, *has a propositional symbol* π_{ψ_i}. *Then a function* $\sigma : Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big) \to Sen^{\mathcal{H}PL}\big((\Delta, Prop)\big)$ *is defined to replace the base sentences that occur in* ρ *and* B_ρ *by propositions from Prop. Formally,*

$$
\begin{aligned}
\sigma(\neg\rho) &= \neg\sigma(\rho)\\
\sigma(\rho \barwedge \rho') &= \sigma(\rho) \barwedge \sigma(\rho')\\
\sigma(i) &= i\\
\sigma(@_i\rho) &= @_i\sigma(\rho)\\
\sigma(\langle\lambda\rangle\rho) &= \langle\lambda\rangle\sigma(\rho)\\
\sigma(\forall x\ \rho) &= \forall x\ \sigma(\rho)\\
\sigma(\mathrm{A}\rho) &= \mathrm{A}\ \sigma(\rho)\\
\sigma(\psi_i) &= \pi_{\psi_i}
\end{aligned}
$$

Definition 14. *For each $\chi \in \Omega_\rho$ we define function $\sigma' : \chi \to Sen^{PL}(Prop)$ such that,*

$$\sigma'(\chi_i) = \begin{cases} \neg\pi_{\psi_i} & \text{if } \chi_i = \neg\psi_i \\ \pi_{\psi_i} & \text{if } \chi_i = \psi_i \end{cases}$$

and denote by $\sigma'[\chi]$ the result of applying σ' to each member of χ.

Note that both σ and σ' are injective.

4.2 Decidability

Lemma 2. *Consider a signature $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$ and $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$. For any $\chi \in \Omega_\rho$, if χ is satisfiable $\sigma'[\chi]$ is also satisfiable.*

Proof. Unsatisfaction of $\sigma'[\chi]$ may only come from the following cases:

1. A component of $\sigma'[\chi]$ is unsatisfiable,
2. two components of $\sigma'[\chi]$ contradict each other.

A component in $\sigma'[\chi]$ is π_{ψ_i} or $\neg\pi_{\psi_i}$, hence the first case never happens. If two elements contradict each other, that is, if one is π_{ψ_i} and the other $\neg\pi_{\psi_i}$ then surely χ has elements ψ_i and $\neg\psi_i$, which renders it unsatisfiable.

Theorem 2. *Consider signature $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$ and $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$. If ρ is satisfiable, $\sigma(\rho)$ is also satisfiable.*

Proof. If ρ is satisfiable we have a model $M = (W, R, m) \in |Mod^{\mathcal{H}I}\big((\Delta, \Sigma)\big)|$ such that $M \models^w \rho$ for some $w \in W$. Through this assumption and Lemma 2, we define a model $(W, R, m') \in |Mod^{\mathcal{H}PL}\big((\Delta, \Sigma)\big)|$ as follows: for any $w \in W$, $m'(w)$ is a model satisfying $\sigma'[\Omega_\rho^{m(w)}]$ (recall that Lemma 2 proves that $\sigma'[\Omega_\rho^{m(w)}]$ is satisfiable).

It remains to show that $(W, R, m') \models^w \sigma(\rho)$, for some $w \in W$. Since models (W, R, m) and (W, R, m') have the same Kripke structure and $\rho, \sigma(\rho)$ only differ in the base sentences, we just need to check that for all $\chi \in \Omega_\rho$, $m(w) \models \chi$ entails that $m'(w) \models \sigma'[\chi]$ for any $w \in W$. Actually, this is a direct consequence of condition, $m(w) \models \Omega_\rho^{m(w)}$ entails that $m'(w) \models \sigma'[\Omega_\rho^{m(w)}]$ for all $w \in W$, which is freely given by the definition of (W, R, m').

Now, we want to show the converse of Theorem 2. For this we need yet another definition to cater for the "preservation" of information with respect to satisfiability of the base sentences; information that is "lost" by $\sigma(\rho)$. Thus,

Definition 15. *Let $Sat^{\mathcal{I}}$ be an effective decision procedure of $\mathcal{I}$, and $\bigvee$ denote the disjunction operator, built from $\mathbb{\wedge}, \neg$. Then define*

$$\eta(\rho) = \begin{cases} \bigvee\{\chi \in \Omega_\rho \mid Sat^{\mathcal{I}}(\chi) \ \textit{is "unsat"}\}, & \textit{if } B_\rho \neq \emptyset \\ \bot, & \textit{otherwise} \end{cases}$$

Corollary 2. *Clearly, satisfiability of ρ entails satisfiability of $\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho)$, which in turn, by Theorem 2, entails satisfiability of $\sigma\big(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho)\big)$.*

Lemma 3. *Consider a model $(W, R, m) \in |Mod^{\mathcal{H}PL}\big((\Delta, Prop)\big)|$ such that $(W, R, m) \models \sigma\big(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho)\big)$. For any $\chi \in \Omega_\rho$, if $\sigma'[\chi]$ is satisfied by a model in $img(m)$, χ is satisfiable.*

Proof. If χ is unsatisfiable then, by definition of η, occurs as one of the literals in $\eta(\rho)$, hence no model in $img(m)$ may satisfy it.

Theorem 3. *Consider signature $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$ and $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$. If $\sigma(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho))$ is satisfiable, then ρ is satisfiable.*

Proof. If $\sigma(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho))$ is satisfiable we have a model $M = (W, R, m) \in |Mod^{\mathcal{H}PL}\big((\Delta, Prop)\big)|$ such that $M \models^w \sigma(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho))$ for some $w \in W$. Through this assumption, and by Lemma 3, we define a model $(W, R, m') \in |Mod^{\mathcal{H}I}\big((\Delta, \Sigma)\big)|$ as follows: for any $w \in W$, $m'(w)$ is a model satisfying χ where $\sigma'[\chi] = \sigma'[\Omega_\rho^{m(w)}]$

It remains to show that $(W, R, m') \models^w \rho$ for some $w \in W$. Since models (W, R, m) and (W, R, m') have the same Kripke structure satisfied by the sentences ρ, $\sigma(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho))$, we just have to show that for all $\chi \in \Omega_\rho$, $m(w) \models \sigma'[\chi]$ entails that $m'(w) \models \chi$ for any $w \in W$. Actually, this is a direct consequence of condition, $m(w) \models \sigma'[\Omega_\rho^{m(w)}]$ entails $m'(w) \models \Omega_\rho^{m(w)}$, for all $w \in W$, which is given by the definition of (W, R, m').

Corollary 3. *From Corollary 2 and Theorem 3 we have that*

$$\rho \textit{ is satisfiable iff } \sigma(\rho \mathbb{\wedge} \mathrm{A}\neg\eta(\rho)) \textit{ is satisfiable.}$$

Then, since $\mathcal{H}'PL$ was already proved to be decidable [12], we may use an effective decision procedure of $\mathcal{H}'PL$ to check for satisfiability of sentences written in $\mathcal{H}'I$. This leads to the expected result

Corollary 4. *If I is decidable then $\mathcal{H}'I$ is also decidable.*

Note that the proof of Theorem 3 paves the way for an example decision algorithm, that is, an algorithm able not only to answer "yes" or "no" to the question "Is ρ satisfiable?", but also to build a model that satisfies sentence ρ, if it exists. Technically, to construct such an algorithm one also needs to have example decision algorithms for both I and $\mathcal{H}'PL$. The latter has at least one prover that meets this requirement [12]. Then, as indicated in the proof, through a $\mathcal{H}'PL$'s decision procedure, one extracts a Kripke frame for ρ in which suitable models of I are "attached" given its example decision algorithm for I.

Finally, note that the decision algorithm for $\mathcal{H}'I$, conceptualised in Theorem 3, may be computationally hard. Indeed, in order to define $\eta(\rho)$ the decision algorithm for I must be executed 2^n times where $n = |B_\rho|$.

In addition, if we want the algorithm to give example models, the decision procedure for I must also be executed a number of times that can be equal the number of worlds in the model built by the decision procedure for $\mathcal{H}'PL$.

4.3 Soundness and completeness

In this section we focus on the entailment system for $\mathcal{H}I$, *i.e.*, on functor $Prf^{\mathcal{H}I}$, to show that the rules in $Prf^{\mathcal{H}I}$ are both sound and complete. Recall that hybridised logics have the typical boolean connectives and that $Prf^{\mathcal{H}I}$ has the *reductio ad absurdum* property. Of course, this means that proving soundness and completeness of $Prf^{\mathcal{H}I}$ boils down to show the equivalence,

$$\rho \text{ is satisfiable iff } \rho \longrightarrow \bot \text{ is not in } Prf^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$$

Recall also that it is assumed that the base institution has the typical boolean connectives and the explicit satisfaction property, as well as that its proof system, Prf^I, is sound, complete and has the *reductio ad absurdum* property.

Theorem 4. *If Prf^I is sound, then $Prf^{\mathcal{H}I}$ is also sound.*

Proof. Consider signature $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$ and $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$.
If $Prf^{\mathcal{H}I}$ is sound then sentence ρ, being satisfiable means that there is no proof arrow $\rho \longrightarrow \bot$ in $Prf^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$. If such an arrow exists, however, it must come from some of the conditions imposed to $Prf^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$, *i.e.*, some of these conditions must be unsound. We check each one:

1. the condition that proof arrows $\emptyset \longrightarrow \rho$ in $Prf^I(\Sigma)$ come to $Prf^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$ is, by assumption, sound.
2. the axioms and proof rules from Table 1 were already proved to be sound (*cf.* [3]).
3. composition, inclusion and product rules are, by definition, sound.

The proof of completeness is more complex. For this we resort to a procedure similar to the one used for proving decidability.

Theorem 5. *Consider a signature $(\Delta, \Sigma) \in |Sign^{\mathcal{H}I}|$ and $\rho \in Sen^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$. If there is no arrow $\rho \longrightarrow \bot$ in $Prf^{\mathcal{H}I}\big((\Delta, \Sigma)\big)$ then there is also no arrow $\sigma(\rho) \longrightarrow \bot$ in $Prf^{\mathcal{H}PL}\big((\Delta, Prop)\big)$,*

Proof. First notice that rules in Table 1 do not distinguish ρ from $\sigma(\rho)$, that is, any such rules may be applied to both sentences. Then observe that, since Table 1 contains all classical tautologies, Prf^{PL} does not bring new rules to $Prf^{\mathcal{H}PL}$ and therefore rules in $Prf^{\mathcal{H}PL}$ are also in $Prf^{\mathcal{H}I}$. Both remarks entail that if there are rules in $Prf^{\mathcal{H}PL}$ that can generate arrow $\sigma(\rho) \longrightarrow \bot$, then the same set of rules (also present in $Prf^{\mathcal{H}I}$) can surely generate it there.

Next, we show the converse of Theorem 5 holds as well. For this we define a function to play a role similar to that played by η in sub-section 4.2.

Definition 16. *Consider a signature* $(\Delta, \Sigma) \in |Sign^{\mathcal{HI}}|$ *and* $\rho \in Sen^{\mathcal{HI}}\big((\Delta, \Sigma)\big)$. *Then define,*

$$\eta'(\rho) = \begin{cases} \bigvee\{\chi \in \Omega_\rho |\ \chi \longrightarrow \bot \text{ is in } Prf^I\}, & \text{if } B_\rho \neq \emptyset \\ \bot, & \text{otherwise} \end{cases}$$

Corollary 5. *Clearly if there is no arrow* $\rho \longrightarrow \bot$ *in* $Prf^{\mathcal{HI}}\big((\Delta, \Sigma)\big)$ *then there is also no arrow* $(\rho \Bbbwedge \mathrm{A}\neg\eta'(\rho)) \longrightarrow \bot$ *in* $Prf^{\mathcal{HI}}\big((\Delta, \Sigma)\big)$.

Lemma 4. *Consider a model* $(W, R, m) \in |Mod^{\mathcal{HPL}}\big((\Delta, Prop)\big)|$ *such that* $(W, R, m) \models \sigma\big(\rho \Bbbwedge \mathrm{A}\neg\eta'(\rho)\big)$. *For any* $\chi \in \Omega_\rho$, *if* $\sigma'[\chi]$ *is satisfied by a model member of* $img(m)$, χ *is satisfiable.*

Proof. If χ is unsatisfiable then, by definition of η' and completeness of Prf^I, occurs as one of the literals in $\eta'(\rho)$, hence no model member of $img(m)$ may satisfy it.

Theorem 6. *If* Prf^I *is complete then* $Prf^{\mathcal{HI}}$ *is also complete.*

Proof. We want to prove that given a signature $(\Delta, \Sigma) \in |Sign^{\mathcal{HI}}|$ and a sentence $\rho \in Sen^{\mathcal{HI}}\big((\Delta, \Sigma)\big)$, if no arrow $\rho \longrightarrow \bot$ exists in $Prf^{\mathcal{HI}}\big((\Delta, \Sigma)\big)$ then ρ is satisfiable.

Hence, let us assume that there is no arrow $\rho \longrightarrow \bot$ in $Prf^{\mathcal{HI}}\big((\Delta, \Sigma)\big)$, which by Corollary 5, entails that there is no proof arrow $\sigma(\rho \Bbbwedge \mathrm{A}\neg\eta'(\rho)) \longrightarrow \bot$ in $Prf^{\mathcal{HPL}}\big((\Delta, Prop)\big)$ and therefore means that $\sigma(\rho \Bbbwedge \mathrm{A}\neg\eta'(\rho))$ is satisfiable. In other words, we have a model $M = (W, R, m) \in |Mod^{\mathcal{HPL}}\big((\Delta, Prop)\big)|$ such that $M \models^w \sigma(\rho \Bbbwedge \mathrm{A}\neg\eta'(\rho))$ for some $w \in W$. Then, by Lemma 4, we are able to define a model $(W, R, m') \in |Mod^{\mathcal{HI}}\big((\Delta, \Sigma)\big)|$, in which, for any $w \in W$, $m'(w)$ is a model for χ where $\sigma'[\chi] = \sigma'[\Omega_\rho^{m(w)}]$.

It remains to show that $(W, R, m') \models^w \rho$ for some $w \in W$. Since models (W, R, m) and (W, R, m') have the same Kripke structure satisfied by sentences ρ and $\sigma(\rho \Bbbwedge \mathrm{A}\neg\eta'(\rho))$, it is enough to show that, for all $\chi \in \Omega_\rho$, $m(w) \models \sigma'[\chi]$ entails that $m'(w) \models \chi$ for any $w \in W$. Actually, this is a direct consequence of the fact that $m(w) \models \sigma'[\Omega_\rho^{m(w)}]$ entails $m'(w) \models \Omega_\rho^{m(w)}$, for all $w \in W$, which comes from the definition of (W, R, m').

5 Conclusions and future work

This paper lays the first steps towards the development of (dedicated) proof tools for hybridised logics, by providing an effective decision algorithm for the satisfiability problem. Additionally the systematic hybridisation of the calculus of a base logic was addressed, and shown to preserve both soundness and completeness.

The next step, from an engineering point of view, is, of course, to develop such a generic, dedicated prover for hybridised logics. A comparison with the strategy of using the parametrised translation to FOL will then be due.

In a similar line of research, lies the development of an alternative decision algorithm, that potentially overcomes the problem detected in the definition of η, which involves calling the decision procedure of the base logic 2^n times, for n the number of base sentences in the sentence under consideration. Such an algorithm may be based on the tableau technique (for instance, the one implemented in [12]) which opens a number of branches as the possible ways to build a model satisfying a given sentence. If the sentence is unsatisfiable then all branches must be closed. If any branch remains open then the decision procedure of the base logic is called to try to close it. Thus, the number of times the decision procedure of the base logic is called is much smaller than in the approach discussed here.

Other results in the literature abstract the combination of logics pattern by considering the "top logic" itself arbitrary. Such is the case of what is called *parametrisation* of logics in [4] by C. Caleiro, A. Sernadas and C. Sernadas. Similarly , the recent method of *importing* logics suggested by J. Rasga, A. Sernadas and C. Sernadas [16] aims at formalising this kind of asymmetric combinations resorting to a graph-theoretic approach. In both cases some decidability and completeness results are given. It should be interesting to see in which ways the hybridisation method relates to these works.

Acknowledgements. This work is funded by ERDF - European Regional Development Fund, through the COMPETE Programme, and by National Funds through FCT within project FCOMP-01-0124-FEDER-028923. M. A. Martins was also supported by the project PEst-OE/MAT/UI4106/2014.

References

1. C. Areces and B. ten Cate. Hybrid logics. In P. Blackburn, F. Wolter, and J. van Benthem, editors, *Handbook of Modal Logics*. Elsevier, 2006.
2. Pedro Baltazar. Probabilization of logics: Completeness and decidability. *Logica Universalis*, 7(4):403–440, 2013.
3. T Braüner. *Proof-Theory of Propositional Hybrid Logic*. Hybrid Logic and its Proof-Theory, 2011.
4. Carlos Caleiro, Cristina Sernadas, and Amílcar Sernadas. Parameterisation of logics. In *WADT*, pages 48–62, 1998.
5. Răzvan Diaconescu. *Institution-independent Model Theory*. Birkhäuser Basel, 2008.
6. Răzvan Diaconescu. Institutional semantics for many-valued logics. *Fuzzy Sets Syst.*, 218:32–52, May 2013.
7. Răzvan Diaconescu and Petros Stefaneas. Ultraproducts and possible worlds semantics in institutions. *Theor. Comput. Sci.*, 379(1-2):210–230, July 2007.
8. José Fiadeiro and Amílcar Sernadas. Structuring theories on consequence. In D. Sannella and A. Tarlecki, editors, *Recent Trends in Data Type Specification*, volume 332 of *Lecture Notes in Computer Science*, pages 44–72. Springer Berlin Heidelberg, 1988.
9. Marcelo Finger and Dov Gabbay. Adding a temporal dimension to a logic system. *Journal of Logic, Language and Information*, 1(3):203–233, 1992.

10. Joseph A. Goguen and Rod M. Burstall. Institutions: abstract model theory for specification and programming. *J. ACM*, 39:95–146, January 1992.
11. Joseph A. Goguen and José Meseguer. Models and equality for logical programming. In Hartmut Ehrig, Robert Kowalski, Giorgio Levi, and Ugo Montanari, editors, *TAPSOFT '87*, volume 250 of *Lecture Notes in Computer Science*, pages 1–22. Springer Berlin Heidelberg, 1987.
12. Guillaume Hoffmann and Carlos Areces. Htab: a terminating tableaux system for hybrid logic. *Electr. Notes Theor. Comput. Sci.*, 231:3–19, 2009.
13. Alexandre Madeira, José M. Faria, Manuel A. Martins, and Luís Soares Barbosa. Hybrid specification of reactive systems: An institutional approach. In G. Barthe, A. Pardo, and G. Schneider, editors, *Software Engineering and Formal Methods (SEFM 2011, Montevideo, Uruguay, November 14-18, 2011)*, volume 7041 of *Lecture Notes in Computer Science*, pages 269–285. Springer, 2011.
14. Manuel A. Martins, Alexandre Madeira, Răzvan Diaconescu, and Luís Soares Barbosa. Hybridization of institutions. In A. Corradini, B. Klin, and C. Cîrstea, editors, *Algebra and Coalgebra in Computer Science (CALCO 2011, Winchester, UK, August 30 - September 2, 2011)*, volume 6859 of *Lecture Notes in Computer Science*, pages 283–297. Springer, 2011.
15. Renato Neves, Alexandre Madeira, Manuel A. Martins, and Luís S. Barbosa. Hybridisation at work. In *CALCO TOOLS*, volume 8089 of *Lecture Notes in Computer Science*. Springer, 2013.
16. J. Rasga, A. Sernadas, and C. Sernadas. Importing logics: Soundness and completeness preservation. *Studia Logica*, 101(1):117–155, 2013.

A coinductive animation of Turing Machines

Alberto Ciaffaglione

Dipartimento di Matematica e Informatica
Università di Udine, Italia
alberto.ciaffaglione@uniud.it

Abstract. We adopt corecursion and coinduction to formalize Turing Machines and their operational semantics in the proof assistant `Coq`. By combining the formal analysis of converging and diverging evaluations, our approach allows us to certify the implementation of the functions computed by concrete Turing Machines. Our effort may be seen as a first step towards the formal development of basic computability theory.

1 Introduction

In this paper we present and discuss a formalization of *Turing Machines* (TMs) and their semantics in the `Coq` implementation of the *Calculus of (Co)Inductive Constructions* ($\mathrm{CC}^{(\mathrm{Co})\mathrm{Ind}}$). Actually, we do not find in the literature much mechanization work dealing with *computability theory*, a foundational, major area of computer science, whereas several other domains have benefited, in recent years, from formal developments carried out within mechanized environments.

As far as we know, the most recent contributions are [10, 2, 1, 13]. Norrish [10] develops a proof of equivalence between the recursive functions and the λ-calculus computational models, and formalizes some computability theory results in the `HOL4` system. Other two works are more related to the present one, as focusing on TMs. Asperti and Ricciotti [1] develop computability theory up to the existence of a universal machine, by carrying out their effort from a perspective oriented to complexity theory in `Matita`. Xu, Zhang and Urban [13] prove the correctness of concrete TMs used to address computability theory, and relate TMs to register machines and recursive functions in `Isabelle/HOL`.

The present work is in fact a departure from the two existing formalizations of TMs, due to the two following reasons. On the one hand, we adopt *corecursion* as definition principle and *coinduction* as proof principle (which are not used by the alternative contributions). On the other hand, inspired by our previous effort on unlimited register machines [2], we encode TMs and their operational semantics from the perspective of *program certification*: *i.e.*, we introduce and justify a methodology for addressing the correctness of concrete TMs.

Actually, TMs form an object system which is challenging in several respects. First, TMs may be completely *unstructured*. Second, the *paper tape*, used by TMs as workspace for computing, is infinite in both directions. Moreover, the evaluation of TMs may give rise to *diverging* computations. Therefore, TMs provide with a typical scenario where the user is required to define and reason about *infinite* objects and concepts. To address formally such an object system, in this paper we settle within *Intuitionistic Type Theory*. In this framework, infinite

structures are managed via *coinductive types*: these, roughly speaking, are collections of elements whose construction requires an infinite numbers of steps. In particular, a handy technique for dealing with corecursive definitions and coinductive proofs in $\mathrm{CC}^{(\mathrm{Co})\mathrm{Ind}}$ was introduced by Coquand [4] and refined by Giménez [6]. Such an approach is particularly appealing, because proofs carried out by coinduction are accommodated as any other infinite, corecursively defined object. This technique is mechanized in the proof assistant `Coq` [11].

The motivations to carry out our formalization of TMs in `Coq` are the following. As it is well-known, traditional papers and textbooks about TMs treat the topic at a more superficial level of detail, and in particular the arguments why particular TMs are correct are often left out. Therefore, the mechanization effort in a proof assistant, besides offering the possibility to discover errors, may typically improve the confidence on the subject (*e.g.*, the correctness proofs for concrete TMs in [13] are acknowledged as the most important contribution). Our work has an educational objective as well: on the one hand we try to illustrate the practice of corecursion and coinduction through suggestive examples; on the other hand we develop our formalization methodology step by step, by justifying it in an analytical way and with ramifications at the proof-theoretical level.

We have used, as starting point for our development, the textbooks by Cutland [5] and by Hopcroft et al. [7]. As an effort towards a broader audience, we display rarely `Coq` code in this paper, but present the encoding at a more abstract level (however, the formalization is available as a web appendix [3]).

Synopsis. In the next section we recall TMs, then in the two following sections we introduce their formalization and illustrate the implementation of coinduction in `Coq`. In the two central Sections 5 and 6 we define a big-step operational semantics for TMs and justify the encoding choices via a small-step semantics, respectively. In the core Section 7 we prove the correctness of three sample TMs, then we state final remarks and discuss related and future work.

2 Turing Machines

Turing Machines (TMs), one among the frameworks proposed to set up a formal characterization of the intuitive ideas of computability and decidability, perform algorithms as carried out by a human agent using paper and pencil. In this work we address *deterministic*, single tape TMs, as introduced by Cutland [5].

Alphabet and tape. TMs operate on a *paper tape*, which is *infinite* in both directions and is divided into single squares along its length. Each square is either blank or contains a symbol from a *finite* set of symbols $s_0, s_1, \ldots, s_n$, named the *alphabet* $\mathcal{A}$ (in fact, the "blank" B is counted as the first symbol s_0).

Specification and computation. At any given time, TMs both scan a single square of the tape (via a *reading/writing head*) and are in one of a *finite* number of *states* $q_1, \ldots, q_m$. Depending on the current state q_i and the symbol being scanned s_h, TMs take *actions*, as indicated by a *specification*[1], *i.e.* a *finite*

[1] As said above, we deal with deterministic TMs, *i.e.*, *non-ambiguous* specifications: for every pair q_i, s_h there is at most one quadruple of the form $\langle q_i, s_h, x, q_j \rangle$.

collection of quadruples $\langle q_i, s_h, x, q_j \rangle$, where $i, j \in [1..m]$, $h \in [0..n]$, $x \in \{R, L\} \cup \mathcal{A}$:

$$\langle q_i, s_h, x, q_j \rangle \triangleq \begin{array}{l} 1)\ \textit{if } x{=}R \textit{ then} \text{ move the head one square to the right} \\ \quad \textit{else if } x{=}L \textit{ then} \text{ move the head one square to the left} \\ \quad \textit{else if } x{=}s_k\ (k{\in}[0..n])\ \textit{then} \text{ replace } s_h \text{ with } s_k \\ 2)\ \text{change the state from } q_i \text{ into } q_j \end{array}$$

When provided with a tape, a specification becomes an *individual TM*, which is capable to perform a *computation*: it keeps carrying out actions by starting from the initial state q_1 and the symbol scanned by the initial position of the head.

Such a computation is said to *converge* if and only if, at some given time, there is no action specified for the current state q_i and the current symbol s_h (that is, there is no quadruple telling what to do next). On the other hand, if this never happens, such a computation is said to *diverge*.

Computable functions. TMs may be regarded as devices for computing numerical *functions*, according to the following conventions. First, a natural number m is represented on a tape by an amount of $m{+}1$ consecutive occurrences of the "tally" symbol 1 (in such a way, the representation of the $0{\in}\mathbb{N}$ is distinguished from the blank tape). Further, a machine M computes the *partial* function $f\colon \mathbb{N} \rightharpoonup \mathbb{N}$ when, for every $a, b{\in}\mathbb{N}$, the computation under M, by starting from its initial state and the leftmost 1 of the a representation, stops with a tape that contains a *total* of b symbols 1 *if and only if* $a{\in}dom(f)$ and $f(a){=}b$ (it is apparent that f is undefined on all inputs a that make the computation diverge). n-ary partial functions $g\colon \mathbb{N}^n \rightharpoonup \mathbb{N}$ are computed in a similar way, where the representations of the n inputs are separated by single blank squares.

In this way, computability theory is developed via TMs, leading to the well-known characterization of the class of *effectively computable* functions.

3 Turing Machines in `Coq`

As described in the previous section, TMs are formed by two components: the specification and the tape, whose content in fact instantiates the former, making it executable. Specifications and tapes actually work together, but are evidently independent of each other from the point of view of the formalization matter.

Our encoding of TMs in `Coq` reflects such an orthogonality: in the present work we are mainly involved in the formal treatment of the tape, which is more problematic and particularly delicate; conversely, we do not pursue the specification-component treatment (*automata* are actually supported by `Coq`'s library), thus keeping that part of the formalization down to a minimum.

Specification and tape. Concerning the *specification* part, we represent states via natural numbers, while alphabet symbols and operations performed by the head are finite collections of elements (we fix the alphabet by adding the "mark" symbol 0 to the "blank" B and the "tally" 1 of previous section). Finally, specifications are finite sequences (*i.e.*, lists) of actions (*i.e.*, quadruples)[2]:

$$\textit{State}: \quad p, q, i \ \in\ \mathbb{N}{=}\{0, 1, 2, \ldots\} \quad \text{state}$$

[2] The middle columns display the metavariables and the datatypes they range over.

Sym	$: a, b, c \quad \in$	$\{B, 1, 0\}$	alphabet symbol
$Head$	$: x \quad \in$	$\{R, L, W(a)\}$	head operation
Act	$: \alpha \quad \in$	$State \times Sym \times State \times Head$	action
$Spec$	$: T, U, V ::=$	$(\iota \mapsto \alpha_\iota)^{\iota \in [0..n]}$ $(n \in \mathbb{N})$	specification

To deal with the *tape*, whose squares are scanned by the head and contain the alphabet symbols, we adopt a pair of *streams* (*a.k.a.* infinite sequences), a datatype borrowed from the `Haskell` community, where is named "zipper":

$HTape$	$: l, r, h, k ::=$	$(\iota \mapsto a_\iota)^{\iota \in [0..\infty]}$	half tape (stream)
$Tape$	$: s, t, u \quad ::=$	$\langle\!\langle\, l, r \,\rangle\!\rangle$	full tape (zipper)

The intended meaning of this encoding is that the second stream ($r = r_0{:}r_1{:}\ldots$) models the infiniteness of the tape towards the right, while the first stream ($l = l_0{:}l_1{:}\ldots$) is infinite towards the left. At any time, the head "$\Downarrow$" will be scrutinizing the first symbol of r, which corresponds physically to:

$$\begin{array}{c} \Downarrow \\ \cdots \mid l_1 \mid l_0 \mid r_0 \mid r_1 \mid \cdots \end{array}$$

This representation allows for a direct access to the content of the tape, an operation which has therefore constant complexity.

Transitions. In Cutland's presentation [5], reported in the previous section, specifications are non-ambiguous lists of actions. To make specifications concretely compute, it is hence necessary, given the current state and tape symbol, to extract from such lists the corresponding head operation and target state.

In our encoding, we delegate the responsibility to guarantee the *determinism* of TMs to a *transition* function tr: $Spec \times State \times Sym \to (State \times Head)$. At the moment (to develop the operational semantics of TMs and its metatheory), we postulate just the *existence* of such a function, without implementing it. Since tr is, in general, *partially* defined (as TMs may either converge or diverge), we assume also the existence of an "halting" output, to handle the convergence:

```
Parameter tr: Spec -> State -> Sym -> (State * Head).
Parameter halt: (State * Head).
```

The motivation of our encoding choice, as said at the beginning of the section, is to keep the formal development, when feasible, as minimal as possible, being the modelling and the management of the tape the focus of our work.

4 Coinduction in `Coq`

The formal treatment of circular, infinite data and relations is supported by `Coq` via the mechanism of *coinductive types*. These are types that have been conceived to provide *finite* representations of infinite structures.

First of all, one may represent concrete, infinite *objects* (*i.e.*, data) as elements of coinductive *sets*[3], which are fully described by a set of *constructors*. From a pure logical point of view, the constructors can be seen as *introduction rules*;

[3] Coinductive sets are coinductive types whose type is the sort `Set`.

these are interpreted coinductively, that is, they are applied infinitely many times, hence the type being defined is inhabited by infinite objects:

$$\frac{a \in Sym \quad h \in HTape}{a{:}h \in HTape}\ (HTape)_\infty$$

In this example we have formalized infinite sequences, *i.e.*, *streams*, of symbols in the alphabet $Sym{=}\{B, 1, 0\}$, the coinductive set $HTape$ which we have introduced in the previous section to model the tape of Turing Machines.

Once a new coinductive type is defined, the system supplies automatically the *destructors*, that is, an extension of the native pattern-matching capability, to *consume* the elements of the type itself. Therefore, coinductive types can also be viewed as the *largest* collection of objects closed *w.r.t.* the destructors. We use here the standard *match* destructor to extract the *head* and *tail* from streams:

$$head(h) \triangleq \text{match } h \text{ with } a{:}k \Rightarrow a \qquad tail(h) \triangleq \text{match } h \text{ with } a{:}k \Rightarrow k$$

However, the destructors *cannot* be used for defining functions by *recursion* on coinductive types, because their elements cannot be consumed down to a constant case. In fact, the natural way to allow self-reference with coinductive types is the *dual* approach of *building* objects that belong to them. Such a goal is fulfilled by defining *corecursive* functions, like, *e.g.*, the following ones:

$$Bs \triangleq B{:}Bs \qquad same(a) \triangleq a{:}same(a) \qquad blink(a,b) \triangleq a{:}b{:}blink(a,b)$$
$$merge(h,k) \triangleq \text{match } h \text{ with } a{:}h' \Rightarrow \text{match } k \text{ with } b{:}k' \Rightarrow a{:}b{:}merge(h',k')$$

Corecursive functions yield infinite objects and may have any type as domain (notice that in the last definition we have applied the *match* destructor to the parameters). To prevent the evaluation of corecursive functions from infinitely looping, their definition must satisfy a *guardedness condition*: every corecursive call has to be guarded by at least one *constructor*, and by *nothing but* constructors[4]. This way of regulating the implementation of corecursion is inspired by *lazy* functional languages, where the constructors do not evaluate their arguments. In fact, corecursive functions are never unfolded in `Coq`, unless their components are explicitly needed, "on demand", by a destruction operation.

Given a coinductive set (such as $HTape$ above), no *proof principle* can be automatically generated by the system: actually, proving properties about infinite objects requires the potential of building *proofs* which are infinite too. What is needed is the design of *ad-hoc* coinductive *predicates*[5] (*i.e.*, relations); these types are in fact inhabited by infinite *proof terms*. The traditional example is *bisimilarity*, that we define on streams and name $\simeq \subseteq HTape \times HTape$:

$$\frac{a \in Sym \quad h,k \in HTape \quad h \simeq k}{a{:}h \simeq a{:}k}\ (\simeq)_\infty$$

[4] Syntactically, the constructors guard the corecursive call "on the left"; this captures the intuition that infinite objects are built via the iteration of a productive step.

[5] Coinductive predicates are coinductive types whose type is the sort `Prop`.

Two streams are bisimilar if we can *observe* that their heads coincide and, recursively, *i.e.*, *coinductively*, their tails are bisimilar. Once this new predicate is defined, the system provides a corresponding proof principle, to carry out proofs about bisimilarity: such a tool, named *guarded induction* principle [4, 6], is particularly appealing in a context where proofs are managed as any other infinite object. In fact, a bisimilarity proof is just an infinite proof term built by corecursion (hence, it must respect the same guardedness constraint that corecursive functions have to). The mechanization of the guarded induction principle provides a handy technique for building proofs inhabiting coinductive predicates; such proofs can be carried out *interactively* through the `cofix` tactic[6]. This tactic allows the user to yield proof terms as *infinitely regressive* proofs, by assuming the thesis as an extra hypothesis and using it later carefully, *i.e.*, provided its application is guarded by constructors. In this way the user is not required to pick out a *bisimulation* beforehand, but may build it incrementally, via tactics.

To illustrate the support provided by the `cofix` tactic, we display below the proof of the property $\forall a, b{\in}Sym.\ merge(same(a), same(b)) \simeq blink(a, b)$, in *natural deduction* style[7]. Mimicking `Coq`'s top-down proof development, first the coinductive hypothesis is assumed among the hypotheses[8]; then, the corecursive functions *same*, *blink* and *merge*, in turn, are unfolded to perform a computation step; finally, the constructor $(\simeq)_\infty$ is applied twice. Hence, the initial goal is reduced to $merge(same(a), same(b)) \simeq blink(a, b)$, *i.e.*, an instance of the coinductive hypothesis. Therefore, the user is eventually allowed to exploit (*i.e.*, discharge) such a hypothesis, whose application is now guarded by the constructor $(\simeq)_\infty$. The application of the coinductive hypothesis in fact completes the proof, and intuitively has the effect of repeating ad infinitum the initial fragment of the proof term, thus realizing the "and so on forever" motto:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{a, b{\in}Sym \quad [merge(same(a), same(b)) \simeq blink(a, b)]_{(1)}}{a, b{\in}Sym \quad a{:}b{:}merge(same(a), same(b)) \simeq a{:}b{:}blink(a, b)}\,(\simeq)_\infty, twice}{a, b{\in}Sym \quad merge(a{:}same(a), b{:}same(b)) \simeq a{:}b{:}blink(a, b)}\,(comp.{:}\ merge)}{a, b{\in}Sym \quad merge(same(a), same(b)) \simeq blink(a, b)}\,(comp.{:}\ same, blink)}{\forall a, b{\in}Sym.\ merge(same(a), same(b)) \simeq blink(a, b)}\,(1), (introduction)$$

5 Operational semantics

As stressed in Sections 2 and 3, the semantics of TMs' specifications is parametric *w.r.t.* tapes: computations, induced by specifications, may either converge or diverge, depending on the supplied tape and the initial position of the head (while the initial state is $1{\in}\mathbb{N}$). In Section 3 we have also chosen an encoding for tapes (via a zipper, made of two streams) such that the position of the head is

[6] A tactic is a command to solve a goal or decompose it into simpler goals.

[7] As usual, local hypotheses are indexed with the rules they are discharged by.

[8] According to Gentzen's notation, we write such an hypothesis (among the leaves of the proof tree) within square brackets, to bear in mind that it can be *discharged*, *i.e.*, cancelled, in the course of a formal proof, as it represents a *local* hypothesis.

implicit within the tape itself. Therefore, the semantics of TMs may be defined by considering *configurations* (T, p, s), where T is a specification, p a state, and $s{=}\langle\!\langle l, r{=}r_0{:}r_1{:}\ldots\rangle\!\rangle$ a tape. Some configurations make actually a computation stop, because there is no action specified by T for the current state p and symbol r_0: these configurations will play the role of the *values* of our semantics. In the following, we will denote with $tr(T, p, s)$ the application of the transition function tr, introduced in Section 3: in particular, we will write $tr(T, p, s){=}\downarrow$ for `(tr T p r0)=halt`, and $tr(T, p, s){=}\langle i, x\rangle$ for `(tr T p r0)=(i,x)`.

In this section we define a *big-step* semantics for TMs, which we will argue to be appropriate and then use in the rest of the paper. The *potential* divergence of computations provides us with a typical scenario which may benefit from the use of *coinductive* specification and proof principles. In fact, a faithful encoding has to reflect the separation between converging and diverging computations, through two different judgments. Hence, we define the *inductive* predicate $b_* \subseteq Spec{\times}Tape{\times}State{\times}Tape{\times}State$ to cope with converging evaluations, and the *coinductive* $b_\infty \subseteq Spec{\times}Tape{\times}State$ to deal with diverging ones.

Definition 1. *(Evaluation) Assume* $T{\in}Spec$, $s{=}\langle\!\langle l{=}l_0{:}l_1{:}\ldots, r{=}r_0{:}r_1{:}\ldots\rangle\!\rangle$ *and* $t{\in}Tape$, $p, q, i{\in}State$. *Then,* b_* *is defined by the following inductive rules:*

$$\frac{tr(T,p,s){=}\downarrow}{b_*(T,s,p,s,p)}\ (stop) \qquad \frac{tr(T,p,s){=}\langle i,R\rangle \quad b_*(T,\langle\!\langle r_0{:}l, tail(r)\rangle\!\rangle, i, t, q)}{b_*(T,s,p,t,q)}\ (right)_*$$

$$\frac{tr(T,p,s){=}\langle i,L\rangle \quad b_*(T,\langle\!\langle tail(l), l_0{:}r\rangle\!\rangle, i, t, q)}{b_*(T,s,p,t,q)}\ (left)_*$$

$$\frac{tr(T,p,s){=}\langle i,W(a)\rangle \quad b_*(T,\langle\!\langle l, a{:}tail(r)\rangle\!\rangle, i, t, q)}{b_*(T,s,p,t,q)}\ (write)_*$$

And b_∞ *is defined by the following rules, (this time) interpreted coinductively*[9]*:*

$$\frac{tr(T,p,s){=}\langle q,R\rangle \quad b_\infty(T,\langle\!\langle r_0{:}l, tail(r)\rangle\!\rangle, q)}{b_\infty(T,s,p)}\ (right)_\infty$$

$$\frac{tr(T,p,s){=}\langle q,L\rangle \quad b_\infty(T,\langle\!\langle tail(l), l_0{:}r\rangle\!\rangle, q)}{b_\infty(T,s,p)}\ (left)_\infty$$

$$\frac{tr(T,p,s){=}\langle q,W(a)\rangle \quad b_\infty(T,\langle\!\langle l, a{:}tail(r)\rangle\!\rangle, q)}{b_\infty(T,s,p)}\ (write)_\infty$$

Notice that in the rules above we write r_0 *and* l_0 *for* $head(r)$ *and* $head(l)$*, respectively (see Section 4 for the definitions of the* head *and* tail *functions).* □

[9] The relation b_∞ is the greatest fixed-point of the above rules, or, equivalently, the conclusions of infinite derivation trees built from such rules.

In our semantics, given a specification T, a tape s and a state p, we capture on the one hand the *progress* of both the head and the states transitions, and on the other hand the *effect* of the operations performed by the head itself.

In detail, the intended meaning of $b_*(T, s, p, t, q)$ is that the computation under the specification T, by proceeding with the tape s and from the state p, *stops* in the state q, by transforming s into t. Conversely, $b_\infty(T, s, p)$ asserts that the computation under T, by proceeding with the tape s and from the state p, *loops*: *i.e.*, there exists a state i (reachable from p) such that, afterwards, the computation comes again at state i after a non-zero, finite number of actions. Therefore, a *final* tape *cannot* exist for b_∞, because the initial tape s is scrutinized (and possibly updated) "ad infinitum" in the course of a diverging computation.

Since TMs are not structured, we have embedded in the big-step semantics one alternative *structuring criterion*, *i.e.*, the number of evaluation steps implicit amount. In fact, we have defined a constant (*i.e.*, non-recursive) rule for b_* (the computation stops because no next action exists) and (co)inductive rules for both b_* and b_∞, to address how moving the head and writing on the tape is carried out, respectively, within a converging computation and a diverging one.

We remark again that the benefit of the zipper encoding of tapes (introduced in Section 3) is that every operation of the head may be carried out via basic functions on streams, whose complexity is minimal and constant.

6 Adequacy

To justify our big-step semantics for TMs, we introduce here an alternative, *small-step* semantics *à la* Leroy [9], and prove that they are equivalent.

We first define a *one-step* reduction concept, to express the three basic actions of TMs (*i.e.*, moving the reading head and writing on the current square). Formally, it is defined as a predicate $\rightarrow \subseteq Spec \times Tape \times State \times Tape \times State$, that we write more suggestively as $(T, s, p) \rightarrow (T, t, q)$. Note (again) that, since TMs are not structured, we do not need to define *contextual* reduction rules.

Afterwards, we can formalize the small-step semantics of interest as reduction sequences: *finite* reductions $\stackrel{*}{\rightarrow}$ are the reflexive and *inductive* transitive closure of $\rightarrow$, while *infinite* reductions $\stackrel{\infty}{\rightarrow}$ its *coinductive* transitive closure. We introduce also *finite positive* reductions $\stackrel{+}{\rightarrow}$, a tool that we will exploit in Section 7.

Definition 2. *(Reduction) Assume* $T \in Spec$, $s = \langle\langle l = l_0 : \ldots, r = r_0 : \ldots \rangle\rangle \in Tape$, *and* $p, q \in State$. *The* one-step *reduction relation* $\rightarrow$ *is defined by the following rules*[10]*:*

$$\frac{tr(T, p, s) = \langle q, R \rangle}{(T, \langle\langle l, r \rangle\rangle, p) \rightarrow (T, \langle\langle r_0{:}l, tail(r) \rangle\rangle, q)} \ (\rightarrow_R)$$

$$\frac{tr(T, p, s) = \langle q, L \rangle}{(T, \langle\langle l, r \rangle\rangle, p) \rightarrow (T, \langle\langle tail(l), l_0{:}r \rangle\rangle, q)} \ (\rightarrow_L)$$

[10] We keep using here the notation for transitions defined in Sections 3 and 5.

$$\frac{tr(T,p,s)=\langle q, W(a)\rangle}{(T,\langle\!\langle l,r\rangle\!\rangle,p)\to(T,\langle\!\langle l,a{:}tail(r)\rangle\!\rangle,q)}\ (\to_W)$$

For $t,u{\in}Tape$, $i{\in}State$, finite *reduction $\overset{*}{\to}$ is defined by induction, via the rules:*

$$\frac{}{(T,s,p)\overset{*}{\to}(T,s,p)}\ (\overset{*}{\to}_0)\quad \frac{(T,s,p)\to(T,u,i)\quad (T,u,i)\overset{*}{\to}(T,t,q)}{(T,s,p)\overset{*}{\to}(T,t,q)}\ (\overset{*}{\to}_+)$$

For $t,u{\in}Tape$, $i{\in}State$, finite positive *reduction $\overset{+}{\to}$ is defined by induction:*

$$\frac{(T,s,p)\to(T,t,q)}{(T,s,p)\overset{+}{\to}(T,t,q)}\ (\overset{+}{\to}_1)\quad \frac{(T,s,p)\to(T,u,i)\quad (T,u,i)\overset{+}{\to}(T,t,q)}{(T,s,p)\overset{+}{\to}(T,t,q)}\ (\overset{+}{\to}_+)$$

And infinite *reduction $\overset{\infty}{\to}$ is defined by the following coinductive rule:*

$$\frac{(T,s,p)\to(T,t,q)\quad (T,t,q)\overset{\infty}{\to}}{(T,s,p)\overset{\infty}{\to}}\ (\overset{\infty}{\to}_\infty)$$

We can prove that evaluation and reduction are equivalent concepts, both in their converging and diverging versions. We remark that our proofs are *constructive*, whereas Leroy [9] had to postulate the "excluded middle" for divergence.

Proposition 1. *(Equivalence) Let be $T{\in}Spec$, $s,t,u{\in}Tape$, and $p,q,i{\in}State$.*

1. *If $(T,s,p)\to(T,u,i)$ and $b_*(T,u,i,t,q)$, then $b_*(T,s,p,t,q)$*
2. *If $(T,s,p)\overset{*}{\to}(T,u,i)$ and $b_*(T,u,i,t,q)$, then $b_*(T,s,p,t,q)$*
3. *$b_*(T,s,p,t,q)$ if and only if $(T,s,p)\overset{*}{\to}(T,t,q)$ and $tr(T,q,t)=\downarrow$*
4. *$b_\infty(T,s,p)$ if and only if $(T,s,p)\overset{\infty}{\to}$*
5. *If $(T,s,p)\overset{+}{\to}(T,u,i)$ and $(T,u,i)\overset{+}{\to}(T,t,q)$, then $(T,s,p)\overset{+}{\to}(T,t,q)$*

Proof. 1) By inversion of the first hypothesis. 2) By structural induction on the derivation of $(T,s,p)\overset{*}{\to}(T,u,i)$, and point 1. 3) Both directions are proved by structural induction on the hypothetical derivation, but the direction ($\Leftarrow$) requires also point 1. 4) Both directions are proved by coinduction and hypothesis inversion. 5) By structural induction on the derivation of $(T,s,p)\overset{+}{\to}(T,u,i)$. □

The above results point out that the proof practice of reduction and evaluation is very similar in `Coq`. In fact, the small-step predicate $\overset{*}{\to}$ is slightly less handy, because, to perform a TM action, the user is required to exhibit the witness tape, besides the target state; obviously, the small-step version lacks the "halting" concept (*i.e.*, $tr(T,q,t)=\downarrow$), which is internalized by the big-step judgment.

Streams vs. lists. We complete this section with a digression about a different encoding for tapes that we pursued in a preliminary phase of our research.

Even if streams are a datatype which captures promptly and naturally the infiniteness of tapes, a formalization approach via (finite) *lists* may also be developed. The choice of lists makes apparent the constraint that, when a computation starts, only a finite number of squares is allowed to contain non-blank

symbols (in this case the empty list is intended to represent an infinite sequence of blanks). We note that the representation of numerical functions in Cutland's setting, that we have adopted at the end of Section 2, respects such a constraint.

Therefore, we proceed by encoding the tape through a pair of lists:

$$\begin{array}{lll} HTape_L : ll, rl ::= (\iota \mapsto a_\iota)^{\iota \in [0..n]} & & \text{half tape (list, } n \in \mathbb{N}) \\ Tape_L \quad : sl, tl ::= \langle\!\langle\, ll, rl \,\rangle\!\rangle & & \text{full tape (list-pair)} \end{array}$$

Afterwards, big-step semantics predicates, playing the role of the ones that deal with streams in Section 5, can be introduced. However, since lists (conversely to streams) might be empty, such predicates must take into consideration this extra pattern and manage it via additional rules. Without going into the full details (for lack of space), we display here the rules for the move-R action[11]:

$$\frac{bL_*(T, \langle\!\langle\, B{:}ll, [\,] \,\rangle\!\rangle, i, t, q)}{bL_*(T, \langle\!\langle\, ll, [\,] \,\rangle\!\rangle, p, t, q)}\ (r_{[\,]})_* \qquad \frac{bL_*(T, \langle\!\langle\, a{:}ll, rl \,\rangle\!\rangle, i, t, q)}{bL_*(T, \langle\!\langle\, ll, a{:}rl \,\rangle\!\rangle, p, t, q)}\ (r_L)_*$$

The inductive convergence predicate $bL_* \subseteq Spec \times Tape_L \times State \times Tape_L \times State$ has the same intended meaning of b_*. The coinductive divergence predicate $bL_\infty \subseteq Spec \times Tape_L \times State$, corresponding to b_∞, is defined analogously.

We remark that alternative formalizations of the predicates bL_* and bL_∞ are feasible, as we could keep just one rule for each action (like for b_* and b_∞) by defining two $head_L$ and $tail_L$ functions from scratch to work on list-tapes. In fact, these functions would behave differently than the ones given by `Coq` library, because in this encoding the empty list represents an infinite sequence of blank symbols. After experimenting with this second solution and experiencing that the complexity of proofs does not change, we have selected the former encoding solution, because in such a way we may delegate to the rules themselves the explicit distinction between the empty list and the non-empty one.

In the end we can prove, via the predicates bL_*, bL_∞, that the semantics with streams may mimic that with lists, and a limited form of the opposite result (in the Proposition below we denote with Bs the stream of blank symbols and with "::" a recursive function that appends a list in front of a stream).

Proposition 2. *(Tape) Let be $T \in Spec$, $ll, rl, ll', rl' \in HTape_L$, and $p, q \in State$.*

1. $bL_*(T, \langle\!\langle\, ll, rl \,\rangle\!\rangle, p, \langle\!\langle\, ll', rl' \,\rangle\!\rangle, q) \Rightarrow b_*(T, \langle\!\langle\, ll{::}Bs, rl{::}Bs \,\rangle\!\rangle, p, \langle\!\langle\, ll'{::}Bs, rl'{::}Bs \,\rangle\!\rangle, q)$
2. $bL_\infty(T, \langle\!\langle\, ll, rl \,\rangle\!\rangle, p)$ *if and only if* $b_\infty(T, \langle\!\langle\, ll{::}Bs, rl{::}Bs \,\rangle\!\rangle, p)$

Proof. 1) By structural induction on the hypothetical derivation. 2) Both the directions are proved by coinduction and hypothesis inversion. □

The difficulty of proving the reverse implication of point 1 above depends on the fact that the representation of the tape through lists is not unique (actually, one may append to any list blank symbols at will). Hence, it seems necessary to introduce an *equivalence* relation on list-tapes, to develop their metatheory.

[11] We omit from both the rules the transition conditions, *i.e.*, the premise $tr(T, p, \langle\!\langle\, ll, [\,] \,\rangle\!\rangle) = \langle i, R \rangle$ from $(r_{[\,]})_*$ and $tr(T, p, \langle\!\langle\, ll, a{:}rl \,\rangle\!\rangle) = \langle i, R \rangle$ from $(r_L)_*$.

For this reason (and because lists demand to double the length of proofs, as the predicates involving them have two constructors for every action), we prefer working with streams. In any case, when addressing concrete TMs, in Section 7, we will guarantee that only a finite number of tape squares is non-blank.

7 Certification

In this section we use the big-step predicates b_* and b_∞, introduced in Section 5 and justified in Section 6, to address the *certification* of the partial functions computed by *individual* TMs. This "algorithmic" approach is significant because it realizes a methodology for the formal development of computability theory.

To make concrete TMs compute we must, first of all, instantiate the *transition* function tr, which we have taken as a parameter of our semantics in Section 3. Simply, we assume that specifications (*i.e.*, lists of quadruples) record the transitions by respecting the order of the source states (*i.e.*, first the transitions for state 1, then those for state 2, and so on). Consequently, we implement the transition function as a recursive function that visits such lists: in particular, it is undefined as soon as it finds a state which is greater than the one it is looking for (or, less efficiently, when it reaches the end of a list)[12]. Concerning the states, we use the 0 to represent the halting state, for which no transition is provided; clearly, also other states may act as halting ones.

Now that our machinery is ready to address the correctness of TMs, we point out that their divergence may be caused by different kinds of behavior. It is easy to manage TMs that diverge by scanning a finite portion of a tape. The interesting case is when TMs scan an infinite area of the tape; this may happen by moving the reading head infinitely either just in one direction or in both directions. In this section we address one example for each pattern of behavior, to convey to the reader the intuition that we can master all of them.

First example: R moves. The first partial function that we work out computes the half of *even* natural numbers, and is not defined on *odd* ones:

$$div2(n) \triangleq \begin{cases} n/2 & \text{if } n \in \mathbb{E} \\ \uparrow & \text{if } n \in \mathbb{O} \end{cases}$$

One algorithm that implements the $div2$ function is conceived as follows. Erase the first "1" (which occurs by definition) and move the head to the right; then try to find pairs of consecutive "1": if this succeeds, erase the second "1" and restart the cycle, otherwise (a single "1" is found) move indefinitely to the right. Such an algorithm can be realized, *e.g.*, by the following specification T:

$$\{\langle 1, 1, W(B), 1\rangle, \langle 1, B, R, 2\rangle, \langle 2, 1, R, 3\rangle, \langle 3, B, R, 3\rangle, \langle 3, 1, W(B), 4\rangle, \langle 4, B, R, 2\rangle\}$$

This implementation of the $div2$ function is then certified through the predicates b_* and b_∞; the computation starts from the state 1 and the following tape[13]:

[12] The specification component of TMs might be also represented through *automata*, for which suitable formalizations can be picked out from `Coq`'s library.

[13] From now on, we will use "$a \mid -$" to represent an infinite amount of "a" symbols.

$$- \mid B \mid \overset{\Downarrow}{1} \mid \underbrace{1 \mid - \mid 1}_{\text{n}} \mid B \mid - \tag{1}$$

which is formalized as $\forall n.\, \langle\!\langle Bs,\, 1{:}ones(n){::}Bs \rangle\!\rangle$, where Bs is the stream of blank symbols, $ones(n)$ a list of n "1" symbols, ":" the *cons* operation on lists, and "::" a function that appends a list in front of a stream.

To fulfill our goal we carry out, via tactics, a top-down formal development that simulates the computation of the TM at hand. First, we perform a write-B and a move-R action from the starting configuration[14] (state 1 and tape (1), that represents the input n), thus reaching the state 2 with the tape:

$$- \mid B \mid \overset{\Downarrow}{\underbrace{1 \mid - \mid 1}_{\text{n}}} \mid B \mid - \tag{2}$$

Proving the *divergence* requires a combination of coinductive and inductive reasoning. The core property is the divergence when proceeding from the state 3 and a right-hand blank tape, a lemma which is proved by coinduction[15]:

$$\dfrac{\dfrac{\dfrac{l{\in}HTape \quad tr(T,3,\langle\!\langle l, B{:}Bs\rangle\!\rangle){=}\langle 3,R\rangle \quad [b_\infty(T,\langle\!\langle B{:}l, Bs\rangle\!\rangle,3)]_{(1)}}{l{\in}HTape \quad b_\infty(T,\langle\!\langle l, B{:}Bs\rangle\!\rangle,3)}\,(right)_\infty}{l{\in}HTape \quad b_\infty(T,\langle\!\langle l, Bs\rangle\!\rangle,3)}\,(computation{:}\ Bs)}{\forall\, l{\in}HTape.\ b_\infty(T,\langle\!\langle l, Bs\rangle\!\rangle,3)}\,(1),(introduction)$$

If n is *odd*, we prove by induction on k that the tape (2) leads to divergence:

$$\forall\, l{\in}HTape.\ b_\infty(T, \langle\!\langle\, l,\, ones(2k{+}1){::}Bs\,\rangle\!\rangle, 2)$$

If $k{=}0$, carry out a move-R and apply the lemma above; if $k{=}h{+}1$, complete a cycle (by erasing the second "1") and conclude via the induction hypothesis.

We address the *convergence* in the complementary scenario (an *even* input n in (2)) by proving the following property, again by induction on k:

$$\forall\, l{\in}HTape.\ b_*(T, \langle\!\langle\, l,\, ones(2k){::}Bs\,\rangle\!\rangle, 2, \langle\!\langle\, repeat(k){::}\, l,\, Bs\,\rangle\!\rangle, 2)$$

where $repeat(k)$ in the final tape stands for a list of k consecutive pairs "B:1". □

Second example: R and L moves. The second sample function that we choose is partially defined on input *pairs*, and may be named "partial minus":

$$pminus(m,n) \triangleq \begin{cases} m-n & \text{if } m \geq n \\ \uparrow & \text{if } m < n \end{cases}$$

To compute it, we devise the following algorithm. First scan the tape towards the right till reaching the B that separates the two inputs; then erase the leftmost "1" from the representation of n and the rightmost "1" from that of m (both the

[14] Given a specification T, a configuration will be a pair $\langle state, tape\rangle$ from now on.

[15] As at the end of Section 4, we display coinductive proofs in natural deduction-style: the coinductive hypothesis is indexed with the rule it is discharged by.

"1s" must occur) by replacing them, respectively, with a mark symbol "0" (on the right, for n) and a B (on the left). The core of the computation is repeating this cycle, which leads to one of two possible situations: if the end of n is reached (*i.e.*, we are scanning the first B on the right of a 0-block), then stop; on the other hand, replacing m with B symbols may cause that the head (looking for "1s") moves indefinitely on the left. The specification is the following:

$$U \triangleq \{\langle 1,1,R,1\rangle, \langle 1,B,R,2\rangle, \langle 2,0,R,2\rangle, \langle 2,1,W(0),3\rangle, \langle 3,0,L,3\rangle,\\ \langle 3,B,L,4\rangle, \langle 4,B,L,4\rangle, \langle 4,1,W(B),5\rangle, \langle 5,B,R,5\rangle, \langle 5,0,R,2\rangle\}$$

The initial part of the formal development (erasing the first pair of "1s", so moving from state 1 to 5) is common to the divergence and convergence cases[16]:

$$-\mid B \mid \underset{\Downarrow}{\underbrace{1 \mid - \mid 1}_{\text{m+1}}} \mid B \mid \underbrace{1 \mid - \mid 1}_{\text{n+1}} \mid B \mid - \stackrel{*}{\Longrightarrow} - \mid B \mid \underbrace{1 \mid - \mid 1}_{\text{m}} \mid B \mid B \mid \underset{\Downarrow}{0} \mid \underbrace{1 \mid - \mid 1}_{\text{n}} \mid B \mid -$$

At this point of the proof, the key pattern to be mastered is shaped as follows:

$$-\mid B \mid \underbrace{1 \mid - \mid 1}_{\text{m}} \mid \underbrace{B \mid - \mid B}_{\text{k+2}} \mid \underset{\Downarrow}{\underbrace{0 \mid - \mid 0}_{\text{k+1}}} \mid \underbrace{1 \mid - \mid 1}_{\text{n}} \mid B \mid - \qquad (3)$$

Starting from this tape and the state 5, we can discriminate between divergence and convergence by distinguishing the case $m{<}n$ from $m{\geq}n$. Notice that we have introduced the variable k to obtain a more general induction hypothesis.

When we come to the state 5 and an instance (for $k{=}1$) of the above tape (3) we prove the *divergence*, under the hypothesis $m{<}n$, by nested induction on n and m. This proof requires auxiliary lemmas, to scan the "0s" and "Bs" blocks (by induction on k) and for assuring the divergence from state 4 with a tape of "Bs" towards the left. One key point is that we can use the predicate b_∞ in a *compositional* way (*i.e.*, we split a divergent computation into a convergent one, which can be proved easily, and another divergent one, which becomes our goal); *e.g.*, we scan a 0-block via the following lemma, proved by induction on k:

$$\forall\, k{\in}\mathbb{N}, \forall\, l, r{\in}HTape.\ b_\infty(U, \langle\!\langle\, blanks(k)::l,\ r\,\rangle\!\rangle, 5) \Rightarrow b_\infty(U, \langle\!\langle\, l,\ blanks(k)::r\,\rangle\!\rangle, 5)$$

Conversely, it is *not* possible to use the predicate b_* in a compositional way to manage the *convergence* scenario. The problem is that b_* requires to exhibit the final tape, but in this case, due to the complexity of the proof, we cannot master it *tout-court* as we have done in Example 1. Therefore, we need an extra tool to accomplish the convergence. Actually, such a tool is provided by the *small-step* predicate $\stackrel{*}{\rightarrow}$: by applying the Proposition 1.2, we may decompose a convergent computation and address separately the intermediate steps. In the end, we carry out the proof from (3), under the hypothesis $m{\geq}n$, by nested induction on n and m and by means of lemmas similar to those used for b_∞. □

[16] Informally, we represent with $\stackrel{*}{\Longrightarrow}$ the effect of a finite number of actions on a tape.

Third example: R and L moves, infinitely. In this example we consider the unary function $f_\emptyset$, undefined on every input, for which we devise an implementation that points out a problem about the mechanization of coinduction.

In fact, our algorithm to compute $f_\emptyset$ is simple: first scan the 1-block towards the right and replace the first blank with a "1"; then move back to the left till reaching the first blank and replace it again with a "1"; then proceed infinitely in the same way. The specification we pick out is minimal:

$$V \triangleq \{\langle 1, 1, R, 1\rangle, \langle 1, B, W(1), 2\rangle, \langle 2, 1, L, 2\rangle, \langle 2, B, W(1), 1\rangle\}$$

The idea beneath the formal divergence proof is nesting a couple of inductions inside the main coinduction; that is, by using the notation introduced in the previous Example 2 to display the modification of the tape, we want to perform the two computations (passing to state 2 and then coming back to state 1):

$$-\mid B \mid \underset{\Downarrow}{1} \underbrace{\mid - \mid 1}_{\text{n+1}} \mid B \mid - \overset{*}{\Longrightarrow} - \mid B \mid \underset{\Downarrow}{1} \underbrace{\mid - \mid 1}_{\text{n+2}} \mid B \mid - \overset{*}{\Longrightarrow} - \mid B \mid \underset{\Downarrow}{1} \underbrace{\mid - \mid 1}_{\text{n+3}} \mid B \mid -$$

It is apparent that, to accommodate this proof, we may assume the coinductive hypothesis for the initial configuration (state 1 and leftmost tape above) and then carry out two finite computations, thus reducing to a configuration (state 1 and rightmost tape) which is an instance of the coinductive hypothesis itself.

Nevertheless, the application of the coinductive hypothesis is *not* allowed by `Coq`, because the whole proof (*i.e.*, the proof term built interactively through tactics, and mainly via `cofix`) is recognized as *non-guarded* by constructors. Essentially, this is caused by the fact the syntactic check does not accept an induction (*i.e.*, a lemma) nested inside the coinductive development[17].

To circumvent the problem, we introduce a new divergence predicate, by taking advantage of the small-step concepts defined and analyzed in Section 6. The idea is very direct: a divergent computation may be characterized as the coinductive transitive closure of the inductive reduction relation $\overset{+}{\to}$.

Definition 3. *(Guarded reduction) Assume $T \in Spec$, $s, t \in Tape$, and $p, q \in State$. The* guarded reduction *relation $\overset{\infty}{\Rightarrow}$ is defined by the following coinductive rule:*

$$\frac{(T, s, p) \overset{+}{\to} (T, t, q) \quad (T, t, q) \overset{\infty}{\Rightarrow}}{(T, s, p) \overset{\infty}{\Rightarrow}} \ (\overset{\infty}{\Rightarrow}_\infty)$$

Proposition 3. *(Infinite reduction) Let be $T \in Spec$, $s \in Tape$, and $p \in State$.*

1. *If $(T, s, p) \overset{\infty}{\Rightarrow}$, then $(T, s, p) \overset{\infty}{\to}$*
2. *$b_\infty(T, s, p)$ if and only if $(T, s, p) \overset{\infty}{\Rightarrow}$*

Proof. 1) By coinduction and hypothesis inversion. 2) ($\Rightarrow$) By coinduction and hypothesis inversion. ($\Leftarrow$) By Proposition 1.4 and point 1. □

[17] See in [8] the proposal of an alternative, *semantic* guardedness checking.

Since the reduction predicate $\overset{\infty}{\Rightarrow}$ is equivalent to b_∞, we adopt the former to carry out our divergence proof. Actually, $\overset{\infty}{\Rightarrow}$ does not suffer from the non-guardedness problem, as it is apparent from the following proof tree[18]:

$$\dfrac{\dfrac{n{\in}\mathbb{N} \quad (V,s,1) \overset{+}{\to} (V,t,1) \quad [(V,t,1) \overset{\infty}{\Rightarrow}]_{(1)}}{n{\in}\mathbb{N} \quad (V,s,1) \overset{\infty}{\Rightarrow}}\ (\overset{\infty}{\Rightarrow}_\infty)}{\forall\, n{\in}\mathbb{N}.\ (V,s,1) \overset{\infty}{\Rightarrow}}\ (1), (\mathit{introduction})$$

The proof of the premise $(V,s,1) \overset{+}{\to} (V,t,1)$ relies on the transitivity of $\overset{+}{\to}$ (Proposition 1.5) and on two auxiliary lemmas, argued by induction on n. □

8 Conclusion

In the present contribution we have formalized TMs and their (big-step and small-step) operational semantics in the `Coq` proof assistant. Our key choices are the encoding of tapes as pairs of *streams* (managed by means of corecursion) and a clear distinction between *converging* computations (modeled via inductive predicates) and *diverging* ones (formalized through coinductive predicates). In the previous, core section we have pointed out the potential of our machinery, by proving the correctness of representative TMs (that is, by certifying the implementation of the partial functions computed by them).

Our encoding provides a completely mechanized management of the transitions (via the `auto` tactic), with the benefit that we may concentrate on the formal treatment of the tape and the logic of proofs. *Divergence* can be proved very often in a compositional way, via the sole big-step coinductive predicate. When "non-guardedness" complications arise (essentially because induction is nested inside coinduction), alternative, equivalent coinductive predicates may be employed, by taking advantage of their close relationship with the small-step semantics concepts. On the other hand, it is not always possible to master *convergence* proofs by compositionality. When this is not feasible (due to the difficulty of the proof at hand), the small-step semantics predicates may be used again as an auxiliary tool, to perform intermediate computation steps.

We note that, in order to carry out either divergence or convergence proofs, often the user has the responsibility to figure out how to decompose the main goal. As usual, it is sometimes necessary to generalize the statements to obtain sufficiently powerful (co)inductive hypotheses. Moreover, some proofs require a subtle combination of inductive and coinductive reasoning.

Related work. The contributions of the literature mainly related to the present one are those by Asperti and Ricciotti in `Matita` [1], Xu, Zhang and Urban in `Isabelle/HOL` [13], and Leroy in `Coq` [9]. Both the first two works address TMs, achieving the ambitious goals we have reported in Section 1.

Asperti and Ricciotti formalize the tape as a triple, made of two lists plus the square currently scrutinized. The non-termination is managed by requiring

[18] We write s for $\langle\!\langle\, Bs,\, ones(n{+}1){::}Bs \,\rangle\!\rangle$ and t for $\langle\!\langle\, Bs,\, ones(n{+}3){::}Bs \,\rangle\!\rangle$.

that the total computation function returns an optional value, when it meets an upper bound of iterations without reaching a final state; the semantics is defined through a relation between tapes, which may be (weakly) realized by TMs. Xu, Zhang and Urban represent the tape via a pair of lists. They handle the non-termination in a similar way, *i.e.*, via the condition that there is no transition into a halting state; the semantics is then defined by means of Hoare-rules.

None of these two works makes use of corecursion and coinduction (that we have exploited to deal with stream-tapes and divergence); from this perspective, our paper is more related to that of Leroy [9], who adopts coinduction in `Coq` to capture infinite evaluations and reductions of a call-by-value λ-calculus.

Future work. We believe that the main result achieved by our work (*i.e.*, the development of a technology for proving the correctness of concrete TMs, via several versions of big-step and small-step semantics) is a promising tool to pursue more advanced goals which are outside the scope of the present paper.

In particular, our effort may be seen as a first step towards the development of computability theory, as the construction of "brick" TMs and their composition at higher-levels of abstraction is the natural progress of this contribution.

Also, it would be stimulating to relate the present formalization to that of unlimited register machines, that we have addressed in a previous work [2].

References

1. A. Asperti and W. Ricciotti. Formalizing Turing Machines. In Proc. of *WoLLIC*, *LNCS* 7456, pp. 1–25. Springer, 2012.
2. A. Ciaffaglione. A coinductive semantics of the Unlimited Register Machine. In Proc. of *INFINITY*, *EPTCS* 73, pp. 49–63, 2011.
3. A. Ciaffaglione. The Web appendix of this paper, 2014. Available at `http://sole.dimi.uniud.it/~alberto.ciaffaglione/Turing/`.
4. T. Coquand. Infinite objects in Type Theory. In Proc. of *TYPES*, *LNCS* 806, pp. 62–78. Springer, 1993.
5. N. J. Cutland. *Computability: An Introduction to Recursive Function Theory.* Cambridge University Press, 1980.
6. E. Giménez. Codifying guarded definitions with recursive schemes. In Proc. of *TYPES*, *LNCS* 996, pp. 39–59. Springer, 1994.
7. J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to automata theory, languages, and computation.* Addison-Wesley, 2003.
8. C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis. The power of parameterization in coinductive proof. In Proc. of *POPL*, pp. 193–206. ACM, 2013.
9. X. Leroy. Coinductive big-step operational semantics. In Proc. of *ESOP*, *LNCS* 3924, pp. 54–68. Springer, 2006.
10. M. Norrish. Mechanised Computability Theory. In Proc. of *ITP*, *LNCS* 6898, pp. 297–311. Springer, 2011.
11. The `Coq` Development Team. *The Coq Proof Assistant, version 8.4.* INRIA, 2012. Available at `http://coq.inria.fr`.
12. A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc.*, 42, 1936.
13. J. Xu, X. Zhang, and C. Urban. Mechanising Turing Machines and Computability Theory in Isabelle/HOL. In Proc. of *ITP*, *LNCS* 7998, pp. 147–162. Springer, 2013.

Mechanised Semantics of BSP Routines with Subgroup Synchronisation

Jean Fortin and Frédéric Gava

Laboratory of Algorithms, Complexity and Logic (LACL), University of Paris-East
{jean.fortin,frederic.gava}@univ-paris-est.fr

Abstract. This paper presents a core language for BSP algorithms with subgroup synchronisation. We give two mechanised semantics for this language using COQ and prove some common properties on the semantics.

1 Introduction

In this paper, we present different semantics for a core Bulk-Synchronous Parallel (BSP) language with subgroup synchronisation. The language is the one of our tool for correctness of BSP programs called BSP-WHY [9]. Subgroups allow to synchronise only a part of the processors and to avoid global barriers of all the processors of the parallel machine. The semantics have been written using the COQ proof assistant and we have proved some common properties on the semantics.

1.1 Motivation

Why use BSP? The **B**ulk-**S**ynchronous **P**arallel (BSP) model is a *bridging model* between abstract execution and concrete parallel systems [20]. Its initial goal is to have portable parallel programs with a scalable performance prediction. Without dealing with low-level details of parallel architectures, the programmer can focus on algorithm design —complexity, correctness, *etc.* It is especially suitable for **H**igh-**P**erformance **C**omputing (HPC).

A wide range of current architectures can be seen as BSP computers. For example shared memory multi-cores machines could be used so that each core only accesses a private sub-part of the shared memory and communications could be performed using a dedicated part of the shared memory. Moreover the synchronisation unit is rarely a hardware entity but rather a software component. Supercomputers, clusters of PCs [2], multi-cores [25] and GPUs [13], *etc.* can be thus considered as BSP computers, *i.e.* these architectures tie in the BSP model.

The **M**essage **P**assing **I**nterface (MPI) [21] is a standard API for communication in distributed applications. There are numerous implementations of this API, both commercial and open source. MPI's goals are high performance, scalability, and portability. MPI is *de facto* the standard for high-performance computing today. And MPI is the main library used for implementing other parallel languages such as BSP libraries. MPI in its entirety does not satisfy the BSP model. However, a large number of MPI programs use only global operations [5]. These can be viewed as BSP programs, if we allow BSP programs to synchronise over a subgroup of processes. Some BSP libraries already allow subgroup synchronisation, such as the Paderborn University BSP library (PUB) [4]. In order to be able to study these kinds of distributed programs, it is thus necessary to extend the semantics [10, 24] which are the basis of formal verification tools.

Why formal semantics? Solving a problem on a parallel machine is often a complex job. High-level tools (models, languages, *etc.*) are necessary to simplify both the design of parallel algorithms and their programming but also to ensure a better safety of the generated applications [19]. To design tools for proofs of correctness of programs or certified compilations or optimisation, a classical step is to provide operational semantics of the language [16]. A recent approach is to use of theorem provers (*e.g.* COQ) for the development of semantics [16] and then formally prove the properties of the language and correctness of programs. This ensures better safety and trust in the generated software.

Why several semantics? What properties? Big-step (natural) semantics are close to the intuition of the language (a formal specification) and are useful to prove properties. Small-step semantics describe more closely the interleaving of the parallel computations, and are useful to prove program transformations.

Several properties are desired when defining semantics of a parallel language. Because some of the semantics allow to prove some properties easier, if we want to have a coherent work, we must ensure that all semantics are equivalent. For example, it is known that confluence is easier to prove using a big-step semantics rather than using a small-step one. Confluence is a property of ordinary sequential programming languages, but remains an objective for HPC since confluence is a powerful aid in debugging and validating programs even if it limits the programmers flexibility to code certain algorithms [3].

In this article, we give both big-step and small-step semantics of a core language, and we define these semantics using the COQ. This allows us to give a mechanically checked proof of the usual properties that are desired for parallel semantics, ensuring a better confidence in our definitions.

Why a core language? In this article, we define a core BSP language with subgroup synchronisation. WHY-ML is the programming language of the WHY tool for the deductive verification of algorithms [7]. BSP-WHY-ML is the parallel counterpart of our own tool called BSP-WHY [9, 8] for BSP algorithms. Defining a core language allows us to focus on the parallel (BSP) aspects of the semantics, and ignore the language specificities of real-world languages such as C, JAVA, *etc.* Subsequently, this work can be extended to a real-world language with a sufficient team. The BSP model has also concrete implementation in the form of libraries for programming languages such as [2, 4, 25]. Additionally, there are existing tools to transform sequential programs from C and Java to WHY-ML [7], so we will be able in the future to extend those tools to transform real-world BSP programs into BSP-WHY-ML, which first needs operational semantics.

1.2 Outline

We start by giving an introduction to the BSP model (Section 2), and defining our core language for BSP programming with subgroup synchronisation. We then define the semantics of the language in Section 3. We show in Section 4 the mechanisation of those semantics in COQ, and give some results that were proved mechanically. Section 5 discusses some related work and finally, Section 6 concludes the paper and gives a brief outlook to future work.

2 A Core Language for BSP with Subgroups

2.1 BSP Programming with Subgroup Synchronisation

The BSP model. A BSP computer is seen as a set of **p** *uniform* processor-memory pairs connected through a *communication network* [2].

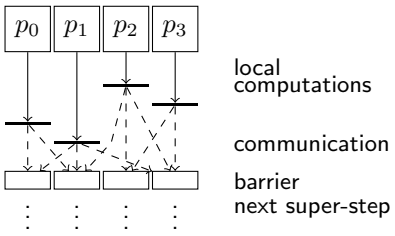

Fig. 1. A BSP super-step.

A BSP program is logically executed as a sequence of *super-steps* (see Fig. 1), each of which is divided into three successive disjoint phases: (1) Each processor only uses its local data to perform sequential computations and to request data transfers to other nodes; (2) The network delivers the requested data; (3) A global synchronisation barrier occurs, making the transferred data available for the next super-step. The BSP model considers communication *en masse*. This is less flexible than asynchronous messages, but easier to debug. Bulk sending also provides better performances since it is faster to send a block of data rather than individual ones. For performance, a BSP library can also send messages during the computation phase, but this is hidden to programmers.

Subgroup synchronisation. The BSP model is based on a global synchronisation. However, in some cases, a parallel algorithm may include problems that can be solved using only a subset of processors. Some libraries, for instance the PUB library [4], extend the basic BSP model, and allow the definition of *subgroups*, which are pairwise disjoint subsets of the set of processors. It is then possible to write a part of the parallel program with the subgroup acting as an independent BSP computer. A call to the `bsp_sync` routine will then synchronise over the subgroup, instead of the whole parallel computer.

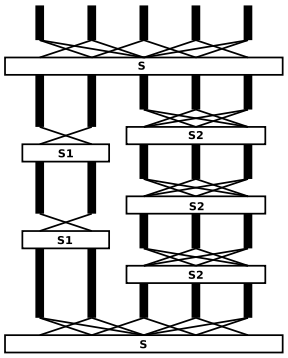

In the left, we show an example of execution of subgroup synchronisation. In this example, the overall group of processors S is split into two subgroups (S1 and S2) which run independent BSP computations. Finally the two subgroups are merged and the whole machine continues its work doing global barriers.

This means that the communication procedures need to be able to tell in which group they are working. This is important for the synchronisation one. An additional argument is thus added to all the parallel procedures, representing a group of processors linked together: a *communicator*. The MPI standard also allows to create communicators: in this way, collective operations are performed only on a subset of the processors, those which participate in the communicator.

2.2 Formal Definition of the Core Language

The syntax of BSP-WHY-ML [8, 9] is the one of WHY-ML [7] with an additional syntax for BSP instructions. Fig. 2 gives the core-calculus. We choose BSP rou-

Pure terms:
t_e ::= $c \mid x \mid \texttt{!}x \mid \phi(t_e, \ldots, t_e)$
Expressions:

e ::=	t_e	term
|	`let` x `=` e `in` e	declaration
|	`let` x `=` `ref` e `in` e	variable
|	`if` e `then` e `else` e	conditional
|	`loop` e	infinite loop
|	`raise` $(E\ e)$	exception
|	`try` e `with` $E\ x \rightarrow e$ `end`	catch it
|	`fun` $x \rightarrow\ e$	pure function
|	x `:=` e	assignment
|	$e\ e$	application
|	`bsp_push` $c\ x$	registering
|	`bsp_pop` $c\ x$	deregistering
|	`bsp_put` $c\ e\ x\ y$	DRMA writing
|	`bsp_get` $c\ e\ x\ y$	DRMA reading
|	`bsp_send` $c\ x\ e$	BSMP sending
|	$\Omega\ e$	Parameters

Fig. 2. Syntax of BSP-WHY-ML.

tines that are close to the ones defined in most BSP libraries: They offer functions for both **B**ulk **M**essage **P**assing (BSMP) and **R**emote **M**emory **A**ccess (DRMA).

Programs contain *pure terms* (t_e, terms without possible side effects) made of constants (integers, booleans, `void`, *etc.*), variables, dereferences (written `!`x), application and application of function symbols ϕ (such as $=$, $\leq$, *etc.*) to pure terms. A special constant *nprocs* (equal to $\mathbf{p}$) and a special variable *pid* (with range $0, \ldots, \mathbf{p}-1$) were also added. In pure terms, we also have introduced the two special function symbols `bsp_nmsg`(t) and `bsp_findmsg` t_1 t_2: the former corresponds to the number of messages received from a processor id t (PUB's "C routine" bsp_nmsgs(t)) and the latter allows to get the t_2-th message from processor t_1 (PUB's "C routine" bsp_findmsg(t1,t2)).

`let`, `if`, `raise`, `try`, `fun` statements are as usual in a ML language. `ref` e introduces a new reference initialized with e that could be modified using `:=`. `loop` e is an infinite loop of body e. The *while* and *sequence* are not part of the core language, but can be easily defined from the *loop* and *let* instructions.

In the core-calculus, the five parallel operations are (where c is for the subgroup): (1) `bsp_push` x, registers a variable x for remote access; (2) `bsp_pop` x, delete x from remote access; (3) `bsp_put` $e\ x\ y$, distant writing of x to y of processor e; (4) `bsp_get` $e\ x\ y$, distant reading from x to y; (5) `bsp_send` $e_1\ e_2$, sending value of e_1 to processor e_2. In order to simplify the presentation of BSP-WHY, parallel operations of the core-calculus (notably DRMA primitives) take simple variables as argument, instead of C's buffers. More details are given in [8].

Ω defines the "parameters" of the semantics, that is routines for which we do not know the code (such as ϕ). It is "user defined": routines can be defined by the user depending on its library or its semantics study. These routines modify the environments (the memory) of the processors. In this way, we abstract the working of some routines: those that can create/delete the subgroups (which are a little different in PUB and MPI) or those that synchronise the processors, that is the routines that perform a barrier (what we call the `SYNC` effect throughout the paper) and enable communication between processors. An example of such a parameter is the standard `bsp_sync`(c) routine of the PUB library, which performs a barrier (synchronises the processors of the subgroup c), and executes all the communications requested during the last superstep by the `put`, `get` and `send` operations (it is painful to define and not really interesting for this work).

The use of Ω also allows to model MPI programs that use only collective operations: routines that define a global communication that involves a group of processes [21]. Previous parallel operations are not modeled using Ω because DRMA operations require a specific treatment with our BSP-WHY tool [8, 9].

3 Formal Operational Semantics

Values to be sent and distant reading/writing are stored in environments of communications as simple list of messages. Aside from the usual environment of variables (here noted $\mathcal{E}$, the usual map from name of variable to values), there are thus six additional components in the environment ($\mathcal{R}$, $\mathcal{C}^{\text{send}}$, $\mathcal{C}^{\text{put}}$, $\mathcal{C}^{\text{get}}$, $\mathcal{C}^{\text{pop}}$, $\mathcal{C}^{\text{push}}$), one per operation that needs communications. Each operation adds a new value to be send in these environments: *e.g.*, a distant writing adds the pair value to be write, to which remote variable, to the $\mathcal{C}^{\text{put}}$ list. We note s the environment of a processor. We note $s.\mathcal{X}$ the access to the component $\mathcal{X}$ of the environment s, $\oplus$ the update of a component of an environment without modifying other components and $\in$ tests the presence of an item in the component.

3.1 Big-step Semantics

Local operations We first define semantics rules for the local execution of a program, on a processor i. We note $s, e \Downarrow^i s', v$ for these local reduction rules (*e.g.* one at each processor i): s is the environment before the execution, e is the program to be executed, s' the environment after the execution and v is the value after execution. It may also be worth mentioning that this being an inductive big-step semantics, the local execution relation is undefined for an infinite loop.

In Fig. 3, we give some examples of the rules for local operations. For each control instruction, it is necessary to give several rules, depending on the result of the execution of the different sub-instructions: one when an execution leads to a synchronisation (when processors finish a super-step, that is `SYNC` effect), one for an exception, and one if it returns directly a value. In the first case, we need to memorise the next instructions of each processor. These intermediate local configurations are noted $\texttt{SYNC}(C, e)$, if e remains to be executed after a synchronisation on the subgroup denoted by C.

With the subgroup synchronisation model, BSP operations are done in the scope of a *communicator*. Every BSP call thus takes an additional argument, this communicator, which describes the subset of processors in which the communications are done. For example, the inductive rule for the "`bsp_send`" primitive manipulates $\mathcal{C}^{\text{send}}$ as follow:

$$\frac{s, e \Downarrow^i s', to \qquad to \in cmt \qquad \{x \mapsto v\} \in s'.\mathcal{E} \qquad s'' = s'.\mathcal{C}^{\text{send}}_{cmt} \oplus \{to, v\}}{s, \texttt{bsp\_send}\ cmt\ x\ e \Downarrow^i s'', \texttt{void}}$$

where "*cmt*" is a valid communicator in the environment. The message to processor "*to*" of the communicator "*cmt*" is added to the queue of messages.

$$\frac{}{s, pid \Downarrow^i s, i} \qquad \frac{}{s, nprocs \Downarrow^i s, \mathbf{p}}$$

$$\frac{s, e_1 \Downarrow^i s', v \qquad s'[x \leftarrow v], e_2 \Downarrow^i s'', o}{s, \texttt{let } x \texttt{ = } e_1 \texttt{ in } e_2 \Downarrow^i s'', o} \qquad \frac{s, e_1 \Downarrow^i s', E(v)}{s, \texttt{let } x \texttt{ = } e_1 \texttt{ in } e_2 \Downarrow^i s', E(v)}$$

$$\frac{s, e_1 \Downarrow^i s', \texttt{SYNC}(C, e')}{s, \texttt{let } x \texttt{ = } e_1 \texttt{ in } e_2 \Downarrow^i s', \texttt{SYNC}(C, \texttt{let } x \texttt{ = } e' \texttt{ in } e_2)} \qquad \frac{s, e \Downarrow^i s', v}{s, x \texttt{ := } e \Downarrow^i s'[x \leftarrow v], void}$$

$$\frac{s, e \Downarrow^i s', E(v)}{s, x \texttt{ := } e \Downarrow^i s', E(v)} \qquad \frac{s, e \Downarrow^i s', \texttt{SYNC}(C, e')}{s, x \texttt{ := } e \Downarrow^i s', \texttt{SYNC}(C, x \texttt{ := } e')} \qquad \frac{s, e\texttt{; loop } e \Downarrow^i s', o}{s, \texttt{loop } e \Downarrow^i s', o}$$

Fig. 3. Big-step semantics: examples of local sequential operations.

Global reduction rules Compared to classical BSP semantics (without subgroup synchronisation), the major changes are located within the parallel rules. Instead of having all the **p** processors synchronise together, and communicate together, it is now possible for a subgroup to synchronize together and make the needed communications. Several subgroups can also work independently from each other, and synchronise at the same time. There are two major options for the semantics in this situation:

1. All processors execute their code locally, until they reach a synchronisation state or they terminate. We execute all possible subgroup synchronisations, and then start again the local computations. We call this option **AllSub**.
2. A subgroup of processors execute their code locally, until they reach their synchronisation. We execute the synchronisation and the associated communications, then start again. We call this option **Diamond**.

With the subgroups, the BSP notion of *superstep* is less clearly defined, and the two formulations could be seen as two different definitions of a superstep in this model. Fig. 4 illustrates them —with one super-step less for **s2** for the **Diamond** rule; because it is too large otherwise. Dotted lines are used when their is no evaluation for a group. In the **Diamond** option we have also give only one possible interleave. In the **AllSub** option, we can notice that the subgroup **s1** does nothing during the third superstep since subgroup **s2** has one more superstep.

There are *pros and cons* to both formulations. In the **AllSub** option, the rule is complex to write, even more so in COQ. However, it is perhaps the rule matching most closely the execution of a parallel program where all processors compute in parallel. The **Diamond** definition is much easier to write and understand, but artificially gives priority to one subgroup over another one. This ordering of the subgroup executions leads to another issue: with this definition, the semantics loses its determinism, since when several subgroup are synchronising, it is possible to choose any subgroup to execute first.

BSP programs are SPMD (**S**ingle **P**rogram **M**ultiple data) ones so an expression e is started **p** times. Different codes can be run by the processors (resp. the subgroups) using conditionals on the "id" of the processors (resp. subgroups). For example "if pid=0 then code1 else code2" for running a different code on processor 0. We model this **p** times executions as a **p**-vector of e with its environment. A

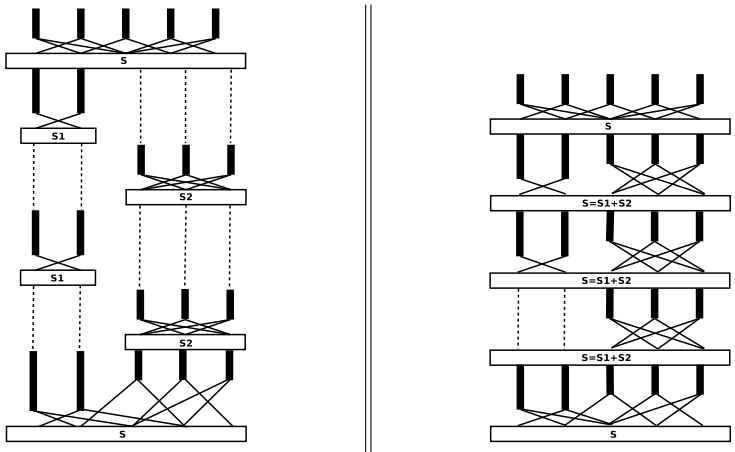

Fig. 4. Illustration of the two formulations for the semantics of subgroups.

final global configuration is a value on all processors that is a **p**-vector of (s_i, v_i), one pair (environment, value) on each processor i. We note $\Downarrow$ for this evaluation.

In Fig. 5, we give 3 variations for the synchronisation rule. The first one is the rule as it would be defined in a semantics without subgroup. The **AllComm** function models the exchanges of messages and thus specifies the order of the received messages depending of the parameter defined in Ω: it modifies the environment of each processor i; it is "just" a reordering of the **p** environments.

The second rule gives the **AllSub** formulation. In this rule, we first partition the set of processors in k subsets that will synchronise, plus a subset N of processors that do not synchronise. **AllCommSub** is then similar to the **AllComm** function of the first rule. However, because there can be several subgroups synchronizing, its exact definition is more complicated. It accepts as argument the set of communicators used in the synchronisation. In addition, it accepts as argument the array of the final values v_i already reached in the super-step. For $i \in N$, the i-th processor terminates without synchronisation with the value v_i, so the i-th component of the result of **AllCommSub** will be the couple (s'_i, v_i). For every set of processors matching a communicator C_j, all the communications corresponding to the communicator are done. Among the messages of these processors, it only considers the ones that were sent within the matching communicator.

Finally, the third rule corresponds to the **Diamond** formulation. **CommDia** is similar to the **AllComm** function, with a few differences. It accepts a second argument (a communicator); it only modifies the environments of the processors in the range of the communicator; and among the messages of these processors, it only considers the ones that were sent within the matching communicator.

It is easy to see that even though this semantics leads to non-determinism, it is still confluent. The reason is that the only source of non-determinism is the communication rules, for which any matching communicator can be chosen. However, at any given point, the communications between two communicators

BSP without subgroup variation:

$$\frac{\forall i \quad s_i, e_i \Downarrow^i s'_i, \texttt{SYNC}(e'_i) \qquad \textbf{AllComm}\{(s'_0, e'_0), \ldots (s'_{\mathbf{p}-1}, e'_{\mathbf{p}-1})\} \Downarrow (s''_0, v_0), \ldots (s''_{\mathbf{p}-1}, v_{\mathbf{p}-1})}{(s_0, e_0), \ldots (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \Downarrow (s''_0, v_0), \ldots (s''_{\mathbf{p}-1}, v_{\mathbf{p}-1})}$$

AllSub variation:

$$\frac{\{0, \ldots, \mathbf{p}-1\} = N \oplus C_1 \oplus \cdots \oplus C_k \qquad \forall i \in C_j \quad s_i, e_i \Downarrow^i s'_i, \texttt{SYNC}(C_j, e'_i) \qquad \forall i \in N \quad s_i, e_i \Downarrow^i s'_i, v_i \qquad \textbf{AllCommSub}\{C_1 \ldots C_k, v, (s'_0, e'_0), \ldots (s'_{\mathbf{p}-1}, e'_{\mathbf{p}-1})\} \Downarrow_{All} (s''_0, v_0), \ldots (s''_{\mathbf{p}-1}, v_{\mathbf{p}-1})}{(s_0, e_0), \ldots (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \Downarrow_{All} (s''_0, v_0), \ldots (s''_{\mathbf{p}-1}, v_{\mathbf{p}-1})}$$

Diamond variation:

$$\frac{\exists C \; \forall i \in C \quad s_i, e_i \Downarrow^i s'_i, \texttt{SYNC}(C, e'_i) \qquad \textbf{CommDia}\{C, (s'_0, e'_0), \ldots (s'_{\mathbf{p}-1}, e'_{\mathbf{p}-1})\} \Downarrow_{Diam} (s''_0, v_0), \ldots (s''_{\mathbf{p}-1}, v_{\mathbf{p}-1})}{(s_0, e_0), \ldots (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \Downarrow_{Diam} (s''_0, v_0), \ldots (s''_{\mathbf{p}-1}, v_{\mathbf{p}-1})}$$

Fig. 5. Big-step semantics: 3 variations for the synchronisation rule.

One (at least) processor diverges:

$$\frac{\exists i \quad s_i, e_i \Downarrow^i_\infty}{\langle (s_0, e_0), \ldots, (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \rangle \Downarrow_\infty}$$

BSP without subgroup variation :

$$\frac{\forall i \quad s_i, e_i \Downarrow^i s'_i, \texttt{SYNC}(e'_i) \qquad \textbf{AllComm}\{\langle (s'_0, e'_0), \ldots, (s'_{\mathbf{p}-1}, e'_{\mathbf{p}-1}) \rangle\} \Downarrow_\infty}{\langle (s_0, e_0), \ldots, (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \rangle \Downarrow_\infty}$$

Diamond variation:

$$\frac{\exists C \; \forall i \in C \quad s_i, e_i \Downarrow^i s'_i, \texttt{SYNC}(C, e'_i) \qquad \textbf{CommDia}\{C, (s'_0, e'_0), \ldots (s'_{\mathbf{p}-1}, e'_{\mathbf{p}-1})\} \Downarrow_\infty}{(s_0, e_0), \ldots (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \Downarrow_\infty}$$

Fig. 6. Global rules for the diverging big-step semantics.

are independent, since each processor leads to a synchronisation in one communicator only. Thus, the diamond property holds.

Co-inductive semantics rules In addition to the standard big-step semantics, it is often useful to define co-inductive (or infinite) semantics rules. They allow to characterize the behaviour of a program that runs indefinitely.

Defining divergence (infinite evaluations) is also done using inference rules but interpreted coinductively. More precisely, the relation is the greatest fixpoint of the rules, or, equivalently, the conclusions of infinite derivation trees built from these rules [17]. Throughout this article, double horizontal lines in inference rules denote inference rules that are to be interpreted coinductively; single horizontal lines denote the inductive interpretation. The co-inductive rules for the local control flow can easily be inferred from the regular big-step semantics rules. We give as an example one of the co-inductive rules for the "if" instruction:

$$\frac{s, e_1 \Downarrow^i s_1, \texttt{true} \qquad s_1, e_2 \Downarrow^i_\infty}{s, \texttt{if } e_1 \texttt{ then } e_2 \texttt{ else } e_3 \Downarrow^i_\infty}$$

The rule can be read as follow: "If e_1 evaluates to *true*, and if e_2 runs infinitely, then the program `if` e_1 `then` e_2 `else` e_3 will run infinitely". Local BSP operation (`push`, `send`, *etc.*) always terminate, so the co-inductive rules can easily be inferred from the regular big-step semantics rules.

On the other hand, and more interestingly, we define the co-inductive global rules in Fig 6. The first rule states that if one of the processors runs infinitely,

then the parallel program will run infinitely. In the other rules, it is said that if all the processors reach a synchronisation barrier, and if the program runs infinitely starting from the resulting state, then the parallel program runs infinitely.

3.2 Small-step Semantics

Small-step semantics specify the execution of a program, one step at a time. A set of rules is repeatedly applied on program states (or configurations), until a final state is reached. If rules can be applied infinitely, it means the program diverges. If at one point in the execution there is no rule to apply, it is a faulty program.

In our parallel case, we will have two kinds of one-step reductions: local ones (on each processor) noted $\stackrel{i}{\rightharpoonup}$ and global ones (for the whole parallel machine) noted $\rightharpoonup$. The whole evaluation $\rightharpoonup^*$ of a program is the transitive and reflexive closure of $\rightharpoonup$. For diverging programs, we note the whole (co-inductive) reduction $\stackrel{\infty}{\Rightarrow}$. All of our semantics are thus a set of "rewriting" rules. As we will see, the small-step semantics is harder to define than the big-step one.

Problems As for the big-step semantics, most of the rules are as usual in programming language semantics. Synchronisation is the only problem. A naive solution would be to define a global rule similar to this:

$$\langle (s_0, \texttt{bsp\_sync}; e_0), \ldots, (s_{\mathbf{p}-1}, \texttt{bsp\_sync}; e_{\mathbf{p}-1}) \rangle \rightharpoonup \langle (s'_0, e_0), \ldots, (s'_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \rangle$$

that is all processors are waiting for a synchronisation and then each processor executes what remains to be done. The problem with this rule is that it cannot evaluate a synchronisation inside a control structure *e.g.* `if` e_1 `then bsp_sync else` e_3. Different solutions exist. First, adding specific global rules for the synchronisation inside each control instruction; the drawback is that this implies too much rules. Second, using a global rule with "contexts" (a context is an expression with a hole, represented with a "_"): the `bsp_sync` instruction replaces the hole within a context on each processor; the drawback is that the use of contexts is not friendly when using a theorem prover such as Coq. Third, in [24] the authors propose the following rule: $s, \texttt{bsp\_sync} \stackrel{i}{\rightharpoonup} s, Wait(skip)$ in addition with rules to propagate this waiting (as the ones of the big-step semantics) and the following rule: $\langle (s_0, Wait(e_0)), \ldots, (s_{\mathbf{p}-1}, Wait(e_{\mathbf{p}-1})) \rangle \rightharpoonup \langle (s_0, e_0), \ldots, (s_{\mathbf{p}-1}, e_{\mathbf{p}-1}) \rangle$. But two subtleties persist: (1) the rules add a *skip* instruction that complicates the proofs; (2) in their work, "$(e_1; \texttt{bsp\_sync}); e_2$" cannot be evaluated, only "$e_1; (\texttt{bsp\_sync}; e_2)$" can. To remedy to the latter problem, in [10] (semantics *without* subgroups), we choose to add the congruence (equivalence) "$(e_1; \texttt{bsp\_sync}); e_2 \equiv e_1; (\texttt{bsp\_sync}; e_2)$" but that also complicates the proofs.

Local rules The solution we propose is the use of a "continuation semantics", in the spirit of the semantics described in [1][1]. This semantics mainly allows a uniform representation of configurations that facilitates the design of lemmas.

[1] Using this semantics we also get for free the evaluation of control structures in C (*e.g.* break and continue in loops) if we want to move to a language such as C.

$$s, nprocs \bullet \kappa \xrightarrow{i} s, \mathbf{p} \bullet \kappa \qquad \| \qquad s, x \texttt{ := } e \bullet \kappa \xrightarrow{i} s, e \bullet (x \texttt{ := \_}) \bullet \kappa$$
$$s, pid \bullet \kappa \xrightarrow{i} s, i \bullet \kappa \qquad \| \qquad s, \texttt{loop } e \bullet \kappa \xrightarrow{i} s, e\texttt{; loop } e \bullet \kappa$$
$$s, v \bullet (\texttt{let } x \texttt{ = \_ in } e_2) \bullet \kappa \xrightarrow{i} s[x \leftarrow v], e_2 \bullet \kappa \qquad \| \qquad s, v \bullet (x \texttt{ := \_}) \bullet \kappa \xrightarrow{i} s[x \leftarrow v], void \bullet \kappa$$
$$s, \texttt{let } x \texttt{ = } e_1 \texttt{ in } e_2 \bullet \kappa \xrightarrow{i} s, e_1 \bullet (\texttt{let } x \texttt{ = \_ in } e_2) \bullet \kappa$$

Fig. 7. Small-step semantics: example of local sequential operations.

$$s, \texttt{bsp\_send } cmt\ x\ e \bullet \kappa \xrightarrow{i} s, e \bullet \texttt{bsp\_send } cmt\ x\ \_ \bullet \kappa$$
$$s, to \bullet \texttt{bsp\_send } cmt\ x\ \_ \bullet \kappa \xrightarrow{i} s', \kappa \quad \text{if } to \in cmt \text{ and } s' = s.\mathcal{C}^{\text{send}}{}_{cmt} \oplus \{to, x\}$$

Fig. 8. Small-step semantics: example of local BSP operation.

$$\frac{s_i, e_i \bullet \kappa_i \xrightarrow{i} s_i', e_i' \bullet \kappa_i'}{\langle (\dots, (s_i, e_i \bullet \kappa_i), \dots) \rangle \rightharpoonup \langle \dots, (s_i', e_i' \bullet \kappa_i'), \dots \rangle}$$

$$\frac{\exists C\ \forall i \in C \qquad O_i \equiv \texttt{bsp\_sync } C \bullet \kappa_i}{\langle (s_0, O_0), \dots, (s_{\mathbf{p}-1}, O_{\mathbf{p}-1}) \rangle \rightharpoonup \mathbf{CommDia}\{C, \langle (s_0, O_0), \dots, (s_{\mathbf{p}-1}, O_{\mathbf{p}-1}) \rangle\}}$$

Fig. 9. Small-step semantics: global reductions.

A configuration is completed with a control stack κ. The final configuration is $(s, void \bullet \epsilon)$, the final environment s with the empty value and with an empty control stack. The control stack represents what has not been executed —where $\bullet$ is an associative operator. There are sequential control operators to handle local control flow. This is close to an abstract machine. In Fig. 7 we give some examples of local rules of control flow and in Fig. 8 for local BSP operations. In the control stack we find expressions with holes. Each hole represents the sub-expression that is currently evaluated and is represented with a "_". Most instructions are dealt with by several rules. Generally, the first rule simply puts the instruction in the continuation stack, and sets the first basic element of the instruction to compute as the main program. Then, one or more rules will match the possible results of this execution, and perform the necessary operations for the control instruction. For instance, with the first rule of the "if" statement, the "if" continuation is put in the control stack and, depending on the result of e_1, the control stack gives the evaluation of e_2 or e_3 with the rest of the stack. A communication primitive consists in simply adding a new value in the environment.

Global rules As previously, global rules are mainly used to call local ones and $\mathbf{p}$ configurations have to be reduced. Fig. 9 gives those rules. First, the global reduction calls a local one. This represents a reduction by a single processor, which thus introduces an interleaving of computations. Communication and BSP synchronisation (`SYNC` effect) are done with the second rule: each processor i of the communicator C is in the case of a synchronous primitive (such as `bsp_sync`) with its control stack κ_i; that is noted O_i. **CommDia** computes the communication, and returns the new environments. Then what remains to be executed is only the control stacks since the synchronisation has been performed.

Co-inductive semantics rules The whole evaluation $\rightharpoonup^*$ is defined as the reflexive and transitive closure of $\rightharpoonup$. Co-inductive semantics rules are much easier to define with the small-step one. A program runs indefinitely if it has an infinite sequence of small-step reductions. The definition of $\rightharpoonup_\infty$ is the same as in [17].

4 Mechanised Semantics in Coq and Properties

4.1 Mechanised BSP Semantics in Coq

In this section, we present the key ideas of our COQ development. The full COQ source files are available at `http://lacl.fr/gava/sbmf2014_coq.tar.gz`. In [8], more details can be found on the relation between the formalization in COQ and the semantics (of the previous section) at a paper-and-pencil level .

Memory model and environment In BSP-WHY, all variables exchanged contain data of the generic type `value`, which can represent any elementary type. We define a corresponding type in COQ. A few special values are also defined, such as `null`, `void`, `true`, `false`, *etc.* The memory is then defined as a function from memory blocks to values. The blocks are numbered according to numbers. We represent identifiers as positive numbers. The link between a variable and its memory block is then stored as a part of the execution environment, which we will detail more in the next section.

A subgroup is defined as a function from the processor identifiers to the booleans. A communicator, however, can not be simply seen as a subgroup, as it is possible that several communicators share the same subgroups of processors. Instead, we define a communicator as a unique positive number, as we did for other identifiers. Environments are defined using record types, following the definition given in Section 3. It is now possible to define our semantics in COQ, according to the rules given in previous sections.

Big-step semantics For the big-step semantics, there are two parts in the semantics. First, we will define the local reduction rules, which represent the evaluation on a single processor, and then we will give the parallel reduction rules.

As is usual in COQ, the semantics rules are given as an inductive predicate — or co-inductive for infinite reductions. For the local reduction rules, `eval_expr i e a e' o` defines the evaluation of the expression `a` in the environment `e` on the processor `i` (which will be globally defined as a COQ's variable):

```
Variable i: pid.
Inductive eval_expr: env → expr → env → outcome → Prop :=
 | eval_Elet : ∀x a1 a2 e e' e'' v o, eval_expr e a1 e' (Outval v) →
                 eval_expr (update e' x v) a2 e'' o → eval_expr e (Elet x a1 a2) e'' o
 | eval_Eraise : ∀a e e' v ex, eval_expr e a e' (Outval v) →
                                   eval_expr e (Eraise ex a) e' (Outexn ex v)
 | eval_Esync : ∀e, eval_expr e (Esync) e (Outsync sync Evoid)
 | ...
```

The definition of an outcome directly follows the definition given in Section 3: there are three possible outcomes, either the computation returns a value `Outval` (case of the "let" rule), raises an exception `Outexn` (case of "raise"), or requests a synchronisation `Outsync` (case of a synchronous operation), with another expression remaining to be executed.

The definition of the `eval_expr` predicate closely matches the definition of the semantics given in Section 3. We thus have several rules for each language instruction, depending on the kind of outcome obtained during the evaluation of the sub-expressions. There is typically one rule for a value outcome, one rule for an exception outcome, one rule for a synchronisation outcome, *etc.* The complete rules are available in the COQ development: We had two parallel rules in the definition of the semantics, so we have two matching rules in our COQ development, similarly for the co-inductive rules for the infinite evaluations.

Small-step semantics The small-step semantics use the same environments and memory model as the big-step semantics.

However, the semantics are significantly different. Since we chose to use continuation semantics, we need to define the continuations. For each statement in the language, zero, one, or several continuations will be defined, depending on the number of computations that are done sequentially in the statement. For instance, a "if" statement has one. The continuation stack is defined inductively, in a similar fashion as a list:

```
Inductive cont : Type :=
  | Kempty: cont
  | Klet: expr→ cont→ cont
  | ...
```

`Kempty` is the empty continuation, and every other continuation is linked to a previous continuation. `Klet` is the continuation for the "let" statement (and there is one continuation case for each different statement of the language). As show in Fig. 7, the "let" small-step rule, we need the second expression e_2 of the "let" following the linked continuation and the hole replace e_1, hence the need of only one `expr` in the definition.

A major difference with the big-step semantics is the notion of program execution "state". For the big-step semantics, a state was simply the association of a program to execute and an environment. In the small-step semantics however, we define four kinds of states: (1) a "normal" state is similar to the notion of state in the big-step semantics; (2) result states are the values returned when a computation is finished; (3) an Error state is characterized by an error type, and a parameter value; (4) and finally, the synchronisation state is the result of a call to a synchronising parameter.

States are thus naturally defined as an inductive type, with four constructors:

```
Inductive state : Type :=
  | State (a : expr) (e : env) (k : cont) : state
  | ResState (v : value) (e : env) (k : cont) : state
  | ErrState (ex : exn) (v : value) (e : env) (k : cont) : state
  | SyncState (c : comm) (e : env) (k : cont) : state.
```

There is always a continuation in a state, but it can be the empty continuation.

A step in the semantics is defined as an inductive from states to states. The definition closely matches the rules that were given in Fig. 7. The parallel

reduction is defined on parallel states (that is, a function from the pid to the local states, and a global environment). From there, the transitive closure `pstar` is defined in a standard manner. For brevity, we do not give the rules here, they are available in the source code.

4.2 Properties and Results

First of all, we defined two alternative rules for the big-step semantics. One is the most direct translation of the BSP model, but is complex to write and to manipulate. The second one, shorter and simpler, relies on the diamond property. We proved in COQ the following two lemmas:

Lemma 1. *The **Diamond** semantics is confluent that is:* $\forall e, s, s'\ s \approx s'$ *if* $(s, e) \Downarrow_{Diam} (s_1, v_1)$ *and* $(s', e) \Downarrow_{Diam} (s_2, v_2)$ *then* $(s_1, v_1) \approx (s_2, v_2)$

Lemma 2. *Both semantics are equivalent that is:* $\forall e, s, s'\ s \approx s'$ *if* $(s, e) \Downarrow_{Diam} (s_1, v_1)$ *and* $(s', e) \Downarrow_{All} (s_2, v_2)$ *then* $(s_1, v_1) \approx (s_2, v_2)$

Where we note (s, e) for the **p** environments and expressions. The lemma justify our choice of privileging the use of the diamond rule. Note that the proofs are done by induction on the derivation of the big-step execution and rely on a novel congruence $\approx$ over the states (values and environments): we need equivalence between **p** processors that can belong to different subgroups. Confluence makes also the hypothesis that all unknown operations of Ω are deterministic.

We were able to prove some results linking the inductive and co-inductive semantics, both for big-step and small-step semantics.

Lemma 3. $\Downarrow$ *and* $\Downarrow_\infty$ *are mutually exclusive that is:* $\forall e, s$ *if* $(s, e) \Downarrow (s_1, v_1)$ *then* $\forall s'\ s \approx s'$ *then* $\neg((s', e) \Downarrow_\infty)$

Lemma 4. $\rightharpoonup^*$ *and* $\rightharpoonup_\infty$ *are mutually exclusive that is* $\forall e, s$ *if* $(s, e) \rightharpoonup^* (s_1, v_1)$ *then* $\forall s'\ s \approx s'$ *then* $\neg((s', e) \rightharpoonup_\infty)$

Co-inductive semantics is also deterministic in the sense that the constructed infinite tree will always be the same. But as our semantics do not currently give the execution traces, this property is not relevant.

The small-step semantics have less rules than the big-step ones. But finding them is much harder and it is thus less a formal specification of the language. It is thus necessary to ensure that it is correct, which we did by proving the equivalence with the big-step semantics.

Lemma 5. $\rightharpoonup^*$ *and* $\Downarrow$ *(resp.* $\rightharpoonup_\infty$ *and* $\Downarrow_\infty$*) are equivalent.*

The proof has some common ground with classic big-step to small-step equivalences, with two difficulties: (1) The small-step semantics allows operations to execute on the different processors in any order, while the big-step semantics fixes the order; (2) The continuations, coupled with the synchronisation, introduce the need for a notion of equivalence between a program and a continuation.

The first implication (big-step to small-step) is done by induction. However, an induction directly on the stated theorem would not be enough, and we need to generalize the result by defining a notion of equivalence between pairs (program, continuation). This is because after a synchronisation, in the big-step semantics we still have a program to execute, while in the small-step semantics it is a continuation. It is then a induction on the derivation of the big-step execution.

For the second implication, we rely of the next lemma. Since the small-step semantics is confluent, we can order the local executions in any order, in particular the order chosen with the big-step semantics (execution on the first processor first until the synchronisation barrier, then the second, *etc.*) The proof is then done by induction on the derivation.

Lemma 6. $\rightharpoonup$ *is confluent.*

The small-step semantics verifies the diamond property: if from one state two steps are possible, they can only be local executions on two different processors. Since the executions are independent within a super-step, we can reach a common state by executing the other computation.

We give for measure of difficulty of the proofs, the number of needed COQ's lines. This is not perfect but easy to count. We compare these numbers to two other developments: (1) the certified C compiler of [15] *without* the memory model; (2) the IMP core language of [17]. We have the following results:

	Language def	Rules	Lemmas:	1	2	3	4	5	6
BSP-WHY-ML	440	696		170	307	65	53	546	270
CompCert	513	1700		1200	no sense	500	undef	1800	undef
IMP	30	60		12	no sense	14	8	53	11

For our case, confluence is easier to prove. Our language definition and the rules contain our own memory model hence a bigger size. Proving that the local (sequential) rules $\Downarrow^i$ are deterministic takes approximately 40 lines. That indicates that the use of the BSP model involves a 4 times increase in the size of the proof. Using the big-step semantics, for exclusivity of finite and infinite *sequential* evaluations, 20 lines are needed: for BSP it is 3 times bigger. For the equivalence of the semantics of sequential programs, 115 lines are needed: 5 times bigger. In conclusion, applying this work to the real-world semantics of CompCert unfortunately seems to be a hard task; It would roughly take 4 times longer and thus need the work of a team bigger than just two researchers.

5 Related Work

To our knowledge, the first work on a formal operational semantic of BSP is [14]: the author gives a small-step semantics using its own primitives of its own core language. Neither mechanised work nor applications have been done. The interests and examples of the use of mechanised semantics for certified program verifiers are given in [16]. In [11], the author gives a mechanised proof of the results of the weakest preconditions calculus used in WHY. A mechanised big-step semantics of WHY is given. The author used massively dependent types whereas we choose a simple model of the language in the spirit of [15].

Different approaches for proofs of BSP programs have been studied. In [10], we presented the correctness of a classical numerical computation using a mechanised operational semantics. In this work, however, we *did not study* subgroup synchronisation which is necessary to take into account MPI's collective operators. The used of continuations also permits to express easily the synchronisation of subgroups of processors. We recall that subgroups are necessary to express many MPI's operators and also allow to write more efficient programs —notable on hierarchical modern architectures, that is clusters of multi-cores. The derivation of imperative BSP programs using the Hoare's axiom semantics has also been studied in [6, 23]. More recently, these works were extended for subgroup synchronisation in [22]. All of these approaches lack of mechanised proofs.

Another work on concurrent threading with barriers is [12]. The authors have developed and proved sound in COQ a concurrent separation logic for barriers of threads. The authors also showcase a program verification toolset that automatically applies the logic rules (Hoare logic) and discharges the associated proof obligations. It is thus a work for derivation of formal specification into correct parallel programs. The drawback is that only programs with a predefined constant number of threads can be considered. For HPC, we prefer to have correct programs for an unknown number of processors.

Studying MPI is challenging due to the number of routines, their concurrent nature and the lack of formal specifications. Even if some works exist [18], some cases are not taken into account because of too much dependence on the architecture. This enormous number of routines in the API makes it difficult to trust any formal specification of a complete MPI. Moreover, a large number of MPI programs use only global operations [5] which have an understandable semantics.

6 Conclusion

In this paper, we defined different operational semantics, in COQ, for a BSP kernel language with subgroup synchronisation. The semantics were proved confluent and equivalent. The big-step semantics used different kinds of value to express the different situations of the program during its execution: exception, true value or synchronisation of a subgroup. The small-step semantics used novel continuations to express more easily the synchronisation mechanism of the BSP model. The proofs were mechanically checked in COQ. The semantics can be used as a basis for verification tools of BSP, as well as for MPI algorithms relying on collective operations. Studying mechanical semantics in COQ of a core language allows to measure the difficulty to move to a real-world parallel language. Confluence is an important property that makes easier code analysis and debugging.

The semantics will be used to prove the transformations of BSP-WHY. Other roads of work include the possibility to prove compilers for BSP programs, as well as the verification of tools translating C or JAVA to BSP-WHY-ML. We will also work on code analysis for optimisation, for example, to find patterns that can used subgroups synchronisation in place of global barriers.

References

1. A. Appel and S. Blazy. Separation Logic for Small-Step Cminor. In *Theorem Proving in Higher Order Logics (TPHOLs)*, volume 4732 of *LNCS*, pages 5–21. 2007.
2. R. H. Bisseling. *Parallel Scientific Computation. A Structured Approach using* BSP *and* MPI. Oxford University Press, 2004.
3. R. L. Bocchino Jr., V. S. Adve, and M. Snir. Parallel Programming Must be Deterministic by Default. In *USENIX, Hot Topics in Parallelism*. 2009.
4. O. Bonorden, B. Judoiink, I. von Otte, and O. Rieping. The Paderborn University BSP (PUB) Library. *Parallel Computing*, 29(2):187–207, 2003.
5. F. Cappello, A. Guermouche, and M. Snir. On Communication Determinism in HPC Applications. In *Comput. Comms. and Networks (ICCCN)*, pages 1–8. IEEE, 2010.
6. Y. Chen and W. Sanders. Top-Down Design of Bulk-Synchronous Parallel Programs. *Parallel Processing Letters*, 13(3):389–400, 2003.
7. J.-C. Filliâtre and C. Marché. The WHY/Krakatoa/Caduceus Platform for Deductive Program Verification. In *Computer Aided Verification (CAV)*, *LNCS*. 2007.
8. J. Fortin. BSP-WHY*: a Tool for Deductive Verification of* BSP *Programs; Machine-checked Semantics and Application*. PhD thesis, University of Paris-East, 2013.
9. J. Fortin and F. Gava. BSP-WHY: An Intermediate Language for Deductive Verification of BSP Programs. In *HLPP*, pages 35–44. ACM, 2010.
10. F. Gava and J.Fortin. Two Formal Semantics of a Subset of the PUB. In *Parallel, Distributed and Network-Based Processing (PDP)*. IEEE Press, 2009.
11. P. Herms. Certification of a Chain for Deductive Verification. COQ *Workshop*. 2010.
12. A. Hobor and C. Gherghina. Barriers in Concurrent Separation Logic: Now With Tool Support! *Logical Methods in Computer Science*, 8(2), 2012.
13. Q. Hou, *et al.* BSGP: BSP GPU Programming. *ACM Trans. Graph.*, 27(3), 2008.
14. D. S. Lecomber. Methods of BSP Programming. PhD thesis, Oxford. 1998.
15. X. Leroy. Formal Verification of a Realistic Compiler. *Comm. of the ACM*, 52(7):107-115, 2009.
16. X. Leroy. Mechanized Semantics: with Applications to Program Proof and Compiler Verification In *Logics and Languages for Reliability and Security*, p. 195-224. 2010.
17. X. Leroy and H. Grall. Coinductive Big-step Operational Semantics. *Inf. Comput.*, 207(2):284–304, 2009.
18. G. Gopalakrishnan, *et al.* Formal Specification of MPI: Case study in Specifying a Practical Concurrent Programming API. *Sci. Comput. Program.*, 76(2):65-81, 2011.
19. Z. Merali. Computational Science: Error, Why Scientific Programming does not Compute. *Nature*, 467(7317):775–777, 2010.
20. D. B. Skillicorn, J. M. D. Hill, and W. F. McColl. Questions and Answers about BSP. *Scientific Programming*, 6(3):249–274, 1997.
21. M. Snir and W. Gropp. MPI *the Complete Reference*. MIT Press, 1998.
22. A. Stewart. A Programming Model for BSP with Partitioned Synchronisation. *Formal Asp. Comput.*, 23(4):421–432, 2011.
23. A. Stewart, M. Clint, and J. Gabarró. Axiomatic Frameworks for Developing BSP-Style Programs. *Parallel Algorithms and Applications*, 14:271–292, 2000.
24. J. Tesson and F. Loulergue. Formal Semantics for the DRMA Programming Style Subset of the BSPLIB Library. In PPAM, n. 4967 of *LNCS*, pages 1122–1129. 2007.
25. A. N. Yzelman and R. H. Bisseling. An Object-oriented BSP Library for Multicore Programming. *Concur. and Comput.: Pract. and Exp.*, 24(5):533–553, 2012.

Formalization of Zsyntax to Reason about Molecular Pathways in HOL4

Sohaib Ahmad[1], Osman Hasan[1], Umair Siddique[1], and Sofiéne Tahar[2]

[1] School of Electrical Engineering and Computer Science (SEECS)
National University of Sciences and Technology (NUST)
Islamabad, Pakistan
{11mseesahmad,osman.hasan,umair.siddique}@seecs.nust.edu.pk
[2] Department of Electrical and Computer Engineering
Concordia University
Montreal, Quebec, Canada
tahar@ece.concordia.ca

Abstract. The behavioral characterization of biological organisms is a fundamental requirement for both the understanding of the physiological properties and potential drug designs. One of the most widely used approaches in this domain is molecular pathways, which offers a systematic way to represent and analyze complex biological systems. Traditionally, such pathways are analyzed using paper-and-pencil based proofs and simulations. However, these methods cannot ascertain accurate analysis, which is a serious drawback for safety-critical applications (e.g., analysis of cancer cells and cerebral malarial network). In order to overcome these limitations, we recently proposed to formally reason about molecular pathways within the sound core of a theorem prover. As a first step towards this direction, we formally expressed three logical operators and four inference rules of Zsyntax , which is a deduction language for molecular pathways. In the current paper, we extend this formalization by verifying a couple of behavioral properties of Zsyntax based deduction using the HOL4 theorem prover. This verification not only ensures the correctness of our formalization of Zsyntax but also facilitates its usage for the formal reasoning about molecular pathways. For illustration purposes, we formally analyze a molecular reaction of the glycolytic pathway leading from D-Glucose to Fructose-1,6-bisphosphate.

1 Introduction

Molecular biology is extensively used to construct models of biological processes in the form of networks or pathways, such as protein-protein interaction networks and signaling pathways. The analysis of these biological networks, usually referred to as biological regulatory networks (BRNs) or gene regulatory networks (GRNs) [10], is based on the principles of molecular biology to understand the dynamics of complex living organisms. Moreover, the analysis of molecular pathways plays a vital role in investigating the treatment of various human infectious

diseases and future drug design targets. For example, the analysis of BRNs has been recently used to predict treatment decisions for sepsis patients [15].

Traditionally, the molecular biology based analysis is carried out by biologists in the form of wet-lab experiments (e.g. [7, 13]). These experiments, despite being very slow and expensive, do not ensure accurate results due to the inability to accurately characterize the complex biological processes in an experimental setting. Other alternatives for deducing molecular reactions include paper-and-pencil proof methods (e.g. using Boolean modeling [27] or kinetic logic [28]) or computer-based techniques (e.g. [29]) for analyzing molecular biology problems. The manual proofs become quite tedious for large systems, where the calculation of unknown parameters takes several hundred proof steps, and are thus prone to human errors. The computer-based methods consist of graph theoretic techniques [21], Petri nets [11] and model checking [3]. These approaches have shown very promising results in many applications of molecular biology (e.g. [8, 14]). However, these methods are not generic and hence have been used to describe some specific areas of molecular biology [4]. Moreover, the inherent state-space explosion problem of model checking [20] limits the scope of this success only to systems where the biological entities can acquire a small set of possible levels.

Theorem proving [12], i.e., a widely used formal methods technique, does not suffer from the state-space explosion problem of model checking, and has also been advocated for conducting molecular biology based analysis [30]. The main idea behind theorem proving is to construct a computer-based mathematical model of the given system and then verify the properties of interest using deductive reasoning. The foremost requirement for conducting the theorem proving based analysis of any system is to formalize the mathematical or logical foundations required to model and analyze that system in an appropriate logic. There have been several attempts to formalize the foundations of molecular biology. For example, the earliest axiomatization even dates back to 1937 [31] and other efforts related to the formalization of biology are presented in [32, 25]. Recent formalizations, based on K-Calculus [6] and π-Calculus [22–24], also include some formal reasoning support for biological systems. But the understanding and utilization of these techniques is very cumbersome for a working biologist as highlighted by Fontana in [9].

In order to develop a biologist friendly formal deduction framework for reasoning about molecular reactions, we propose to formalize the Zsyntax [4] language in higher-order logic. Zsyntax is a formal language that supports modeling and logical deductions about any biological process. The main strength of Zsyntax is its biologist-centered nature as its operators and inference rules have been designed in such a way that they are understandable by the biologists. Traditionally, logical deductions about biological processes, expressed in Zsyntax , were done manually based on the paper-and-pencil based approach. This limits the usage of Zsyntax to smaller problems and also makes the deduction process error-prone due to the human involvement. As a first step towards overcoming this limitation, we formalized the logical operators and inference rules of Zsyntax in higher-order logic [2]. In the current paper, we build upon these for-

mal definitions to verify a couple of key behavioral properties of Zsyntax based molecular pathways using the HOL4 theorem prover. The formal verification of these properties raises the confidence level in our definitions of Zsyntax operators and inference rules, which have complex interrelationships. Moreover, these formally verified properties can be used to facilitate the formal reasoning about chemical reactions at the molecular level. In order to illustrate the usefulness and effectiveness of our formalization for analyzing real-world problems in molecular biology, we present the formal analysis of a molecular reaction of the glycolytic pathway leading from D-Glucose to Fructose-1,6-bisphosphate [4].

Our current framework handles static reactions but it can be further extended to study the reaction kinetics [4] due to the flexibility of Zsyntax . The main motivation behind using higher-order-logic theorem proving in our work is to be able to leverage upon the high expressiveness of higher-order logic and thus reason about differential equations and probabilistic properties, which form an integral part of reaction kinetics. However, the scope of the current paper is on the formalization of Zsyntax based deduction calculus for molecular pathways but this formalization can later be extended to support reaction kinetics as well because it is done in a higher-order-logic theorem prover.

The rest of the paper is organized as follows: Section 2 provides an introduction to Zsyntax and the HOL4 theorem prover. The higher-order-logic formalization of Zsyntax operators and inference rules using HOL4 is described in Section 3. This is followed by the descriptions of the behavioral properties of Zsyntax along with their formal proof sketches in Section 4. The illustrative case study on the glycolytic pathway is presented in Section 4. We conclude the paper in Section 5 while highlighting some interesting potential applications of our work.

2 Preliminaries

2.1 Zsyntax

Zsyntax [4] exploits the analogy between biological processes and logical deduction. Some of the key features of Zsyntax are: 1) the ability to express molecular reactions in a mathematical way; 2) heuristic nature, i.e., if the conclusion of a reaction is known, then one can deduce the missing data from the initialization data; 3) computer implementable semantics. Zsyntax consists of the following three operators:

Z-Interaction: The interaction of two molecules is expressed by the Z-Interaction ($*$) operator. In biological reactions, Z-interaction is not associative.

Z-Conjunction: The aggregate of same or different molecules (not necessarily interacting with each other) is formed using the Z-Conjunction (&) operator. Z-Conjunction is fully associative.

Z-Conditional: A path from A to B under the condition C is expressed using the Z-Conditional ($\rightarrow$) operator as: $A \rightarrow B$ if there is a C that allows it.

Zsyntax supports four inference rules, given in Table 1, that play a vital role in deducing the outcomes of biological reactions:

Table 1. Zsyntax Inference Rules

Inference Rules	**Definition**
Elimination of Z-conditional(→E)	`if C ⊢ (A → B) and (D ⊢ A) then (C&D ⊢ B)`
Introduction of Z-conditional(→I)	`C&A ⊢ B then C ⊢ (A → B)`
Elimination of Z-conjunction(&E)	`C ⊢ (A&B) then (C ⊢ A) and (C ⊢ B)`
Introduction of Z-conjunction(&I)	`(C ⊢ A) and (D ⊢ B) then (C&D) ⊢ (A&B)`

Besides the regular formulas that can be derived based on the above mentioned operators and inference rule, Zsyntax also makes use of *Empirically Valid Formulae* (EVF). These EVFs basically represent the non-logical axioms of molecular biology and are assumed to be validated empirically in the lab.

It has been shown that any biological reaction can be mapped and their final outcomes can be derived using the above mentioned three operators and four inference rules [4]. For example, consider a scenario in which three molecules A, B and C react with each other to yield another molecule Z. This can be represented as a Zsyntax theorem as follows:

$$\texttt{A \& B \& C} \vdash \texttt{Z}$$

The Z-Conjunction operator & is used to represent the given aggregate of molecules and then the inference rules from Table 1 are applied on these molecules along with some EVFs (chemical reactions verified in laboratories) to obtain the final product Z. For the above example, these EVFs could be:

$$\texttt{A * B} \rightarrow \texttt{X} \text{ and } \texttt{X * C} \rightarrow \texttt{Z}$$

meaning that A will react with B to yield X and X in return will react with C to yield the final product Z.

The main contribution of our paper is the formal verification of the Zsyntax based deduction method based on the higher-order-logic formalization of the above-mentioned operators and inference rules using the HOL4 theorem prover. This work will in turn facilitate the derivation of biological reactions within the sound core of HOL4.

2.2 HOL4 Theorem Prover

HOL4 is an interactive theorem prover developed at the University of Cambridge, UK, for conducting proofs in higher-order logic. It utilizes the simple type theory of Church [5] along with Hindley-Milner polymorphism [17] to implement higher-order logic. HOL4 has been successfully used as a verification framework for both software and hardware as well as a platform for the formalization of pure mathematics.

In order to ensure secure theorem proving, the logic in the HOL4 system is represented in the strongly-typed functional programming language ML [19]. An ML abstract data type is used to represent higher-order logic theorems and the

only way to interact with the theorem prover is by executing ML procedures that operate on values of these data types. The HOL4 core consists of only 5 basic axioms and 8 primitive inference rules, which are implemented as ML functions. Soundness is assured as every new theorem must be verified by applying these basic axioms and primitive inference rules or any other previously verified theorems/inference rules.

A HOL4 theory is a collection of valid HOL4 types, constants, axioms and theorems, and is usually stored as a file in computers. Users can reload a HOL4 theory in the HOL4 system and utilize the corresponding definitions and theorems right away. Various mathematical concepts have been formalized and saved as HOL4 theories by the HOL4 users. We utilize the HOL4 theories of Booleans, arithmetics and lists extensively in our work. Table 2 provides the mathematical interpretations of some HOL4 symbols and functions frequently used in this paper.

Table 2. HOL4 Symbols and Functions

HOL Symbol	**Standard Symbol**	**Meaning**
$\wedge$	*and*	Logical *and*
$\vee$	*or*	Logical *or*
$\neg$	*not*	Logical *negation*
::	*cons*	Adds a new element to a list
++	*append*	Joins two lists together
`HD L`	*head*	Head element of list L
`TL L`	*tail*	Tail of list L
`EL n L`	*element*	n^{th} element of list L
`MEM a L`	*member*	True if a is a member of list L
`LENGTH L`	*length*	Length of list L
`FST`	fst (a, b) = a	First component of a pair
`SND`	snd (a, b) = b	Second component of a pair
`SUC n`	$n + 1$	Successor of a *num*

3 Formalization of Zsyntax

We modeled the molecules as variables of arbitrary data types (α) in our formalization of Zsyntax [2]. A list of molecules (α *list*) represents the Z-Interaction or a molecular reaction among the elements of the list. The Z-Conjunction operator forms a collection of non-reacting molecules and can now be formalized as a list of list of molecules (α *list list*). This data type allows us to apply the Z-Conjunction operator between individual molecules (a list with a single element) or multiple interacting molecules (a list with multiple elements). The Z-Conditional operator is used to update the status of molecules, i.e., generate

a new set of molecules based on the available EVFs (wet-lab verified reactions). Each EVF is modeled in our formalization as a pair (α *list* # α *list list*) where the first element is a list of molecules (α *list*) indicating the reacting molecules and the second element is a list of list of molecules (α *list list*) indicating the resulting set of molecules after the reaction between the molecules of the first element of the pair has taken place. A collection of EVFs is represented as a list of EVFs ((α *list* # α *list list*)*list*) in our formalization.

The elimination of Z-Conditional rule is the same as the elimination of implication rule (Modus Ponens) in propositional logic and thus it can be directly handled by the HOL4 simplification and rewriting rules. Similarly, the introduction of Z-Conditional rule can also be inferred from the rules of propositional logic and can be handled by the HOL4 system without the introduction of a new inference rule. The elimination of the Z-Conjunction rule allows us to infer the presence of a single molecule from an aggregate of inferred molecules. This rule is usually applied at the end of the reaction to check if the desired molecule has been obtained. Based on our data types, described above, this rule can be formalized in HOL4 by returning a particular molecule from a list of molecules:

Definition 1. Elimination of Z-Conjunction Rule

```
⊢ ∀ L m. z_conj_elim L m = if MEM m L then [m] else L
```

The function `z_conj_elim` has the data type (α *list* $\rightarrow$ α $\rightarrow$ α *list*). The above function returns the given element as a single element in a list if it is a member of the given list. Otherwise, it returns the argument list as it is.

The introduction of Z-Conjunction rule along with Z-Interaction allows us to perform a reaction between any of the available molecules during the experiment. Based on our data types, this rule is equivalent to the append operation of lists.

Definition 2. Intro of Z-Conjunction and Z-Interaction

```
⊢ ∀ L m n. z_conj_int L m n = FLAT [EL m L; EL n L]::L
```

The above definition has the data type (α *list list* $\rightarrow$ *num* $\rightarrow$ *num* $\rightarrow$ α *list list*). The HOL4 functions `FLAT` and `EL` are used to flatten a list of list to a single list and return a particular element of a list, respectively. Thus, the function `z_conj_int` takes a list `L` and appends the list of two of its elements `m` and `n` on its head.

Based on the laws of stoichiometry [4], the reacting molecules using the Z-Conjunction operator have to be deleted from the aggregate of molecules. The following function represents this behavior in our formalization:

Definition 3. Reactants Deletion

```
⊢ ∀ L m n. z_del L m n = if m > n
                then del (del L m) n
                else del (del L n) m
```

Here the function `del L m` deletes the element at index `m` of the list L and returns the updated list as follows:

Definition 4. Element Deletion

```
⊢ ∀ L. del L 0 = TL L ∧
  ∀ L n. del L (n + 1) = HD L::del (TL L) n
```

Thus, the function `z_del L m n` deletes the m^{th} and n^{th} elements of the given list `L`. We delete the higher indexed element before the lower one in order to make sure that the first element deletion does not effect the index of the second element that is required to be deleted. The above data types and definitions can be used to formalize any molecular pathway (which is expressible using Zsyntax) and reason about its correctness within the sound core of the HOL4 theorem prover.

Our main objective is to develop a framework that accepts a list of initial molecules and possible EVFs and allows the user to formally deduce the final outcomes of the corresponding biological experiment. In this regard, we first develop a function that compares a particular combination of molecules with all the EVFs and upon finding a match introduces the newly formed molecule in the initial list and deletes the consumed instances.

Definition 5. EVF Matching

```
⊢ ∀ L E m n.
z_EVF L E 0 m n =
    if FST (EL 0 E) = HD L
      then (T,z_del (TL L ++ SND (EL 0 E)) m n
      else (F,TL L) ∧
  ∀ L E p m n.
z_EVF L E (p + 1) m n =
    if FST (EL (p + 1) E) = HD L
      then (T,z_del (TL ++ SND (EL (p + 1) E)) m n
      else z_EVF L E p m n
```

The data type of the function `z_EVF` is: (α *list list* $\rightarrow$ (α *list*#α *list list*) *list* $\rightarrow$ *num* $\rightarrow$ *num* $\rightarrow$ *num* $\rightarrow$ *bool* # α *list list*). The function `LENGTH` returns the length of a list. The function `z_EVF` takes a list of molecules `L` and recursively checks its head, or the top most element, against all elements of the EVF list `E`. If there is no match, then the function returns a pair with its first element being false (`F`), indicating that no match occurred, and the second element equals the tail of the input list `L`. Otherwise, if a match is found then the function replaces the head of list `L` with the second element of the EVF pair and deletes the matched elements from the initial list as these elements have already been consumed. This modified list is then returned along with a true (`T`) value, which acts as a flag to indicate an element replacement.

Next, in order to deduce the final outcome of the experiment, we have to call the function `z_EVF` recursively by placing all the possible combinations of the given molecules at the head of list `L` one by one.

Definition 6. Recursive Function for calling `z_EVF`

```
⊢ ∀ L E m n. z_deduction_recur L E m n 0 = (T,L) ∧
  ∀ L E m n q. z_deduction_recur L E m n (q + 1) =
```

```
if FST (z_recur2 L E m n) ⇔ T
  then z_deduction_recur (SND (z_recur2 L E m n)) E
                         (LENGTH (SND (z_recur2 L E m n)) - 1)
                         (LENGTH (SND (z_recur2 L E m n)) - 1) q
  else (T,SND (z_recur2 L E (LENGTH L - 1) (LENGTH L - 1)))
```

The data type of function `z_deduction_recur` is (α *list list* $\rightarrow$ (α *list* #α *list list*) *list* $\rightarrow$ *num* $\rightarrow$ *num* $\rightarrow$ *num* $\rightarrow$ *bool* # α *list list*). It accepts the list of molecules `L` and the list of EVFs `E` along with their corresponding indices `m` and `n`, respectively, and a recursion variable `q`. It returns a pair with the first element being a Boolean flag, which becomes true when there are no more remaining reactions, and the second element being the list of molecules representing the post-reaction state. The function `z_decuction_recur` recursively calls the function `z_EVF` for all possible molecule combinations using the function `z_recur2`, which in turn uses the function `z_recur1` for this purpose. The arguments `m` and `n` of functions `z_recur1` and `z_recur2` are initialized with `LENGTH L` and the sole purpose of these functions is to exhaust all possible combinations of the variables `m` and `n` for the function `z_conj_int`, given in Definition 5. The formalization of the above mentioned functions and more details about their behavior can be obtained from [1, 2].

In order to model a complete experiment for a given list of molecules, the variable of recursion in the function `z_deduction_recur` should be assigned a value that is greater than the total number of EVFs so that the application of none of the EVF is missed. Similarly, the variables `m` and `n` of the function `z_deduction_recur` should be assigned the values of `(LENGTH L - 1)` to ensure that all combinations of the list `L` are checked against the elements of the list of EVFs. Thus, the final deduction function for Zsyntax can be expressed in HOL4 as follows:

Definition 7. Final Deduction Function for Zsyntax

```
⊢ ∀ L E. z_deduction L E =
     SND (z_deduction_recur L E (LENGTH L - 1) (LENGTH L - 1) LENGTH E)
```

The data type of function `z_deduction` is (α *list list* $\rightarrow$ (α *list* # α *list list*) *list* $\rightarrow$ α *list list*). It accepts the initial list of molecules and the list of valid EVFs and returns a list of final outcomes of the experiment under the given conditions, by calling the function `z_decuction_recur`.

The formal definitions, presented in this section, allow us to recursively check all the possible combinations of the initial molecules against the first elements of given EVFs. In case of a match, the corresponding EVF is applied by replacing the reacting molecules with their outcome in the molecule list and the process restarts again to find other possible matches from the new list of molecules. This process terminates when no more molecules are found to be reacting with each other and at this point we will have the list of post-reaction molecules. The desired result can then be obtained from these molecules using the elimination of Z-Conjunction rule, given in Definition 1. The main benefit of the development, presented in this section, is that it facilitates automated reasoning about the molecular biological experiments within the sound core of a theorem prover.

4 Formal Verification of Zsyntax Properties

In order to ensure the correctness and soundness of our definitions, we use them to verify a couple of properties representing the most important characteristics of molecular reactions. The first property deals with the case when there is no combination of reacting molecules in the list of molecules and in this case we verify that after the Zsyntax based experiment execution both the pre and post-experiment lists of molecules are the same. The second property captures the behavior of the scenario when the given list of molecules contains only one set of reacting molecules and in this case we verify that after the Zsyntax based experiment execution the post-experiment list of molecules contains the product of the reacting molecules minus its reactants along with the remaining molecules provided initially. We represent these scenarios as formally specified properties in higher-order logic using our formal definitions, given in the previous section. These properties are then formally verified in HOL4.

4.1 Scenario 1: No Reaction

We verify the following theorem for the first scenario:

Theorem 1.

```
⊢ ∀ E L.
        ∼(NULL E) ∧ ∼(NULL L) ∧
        (∀ a m n. MEM a E ∧ m < LENGTH L ∧ n < LENGTH L
              ⇒ ∼MEM (FST a) [HD (z_conj_int L m n)])
        ⇒ z_deduction L E = L
```

The variables `E` and `L` represent the lists of EVFs and molecules, respectively. The first two assumptions ensure that both of these lists have to be non-empty, which are the pre-conditions for a molecular reaction to take place. The next conjunct in the assumption list of Theorem 1 represents the formalization of the no-reaction-possibility condition as according to this condition no first element of any pair in the list of EVFs `E` is a member of the head of the list formed by the function `z_conj_int`, which picks the elements corresponding to the two given indices (that range over the complete length of the list of molecules `L`) and appends them as a flattened single element on the given list `L`. This constraint is quantified for all variables `a`, `m` and `n` and thus ensures that no combination of molecules in the list `L` matches any one of the first elements of the EVF list `E`. Thus, under this constraint, no reaction can take place for the given lists `L` and `E`. The conclusion of Theorem 1 represents the scenario that the output of our formalization of Zsyntax based reaction would not make any change in the given molecule list `L` and thus verifies that under the no-reaction-possibility condition our formalization also did not update the molecule list.

The verification of this theorem is interactively done by ensuring the no-update scenario for all molecule manipulation functions, i.e., `z_EVF`, `z_recur1`, `z_recur2` and `z_deduction_recur`, under the no-reaction-possibility condition [1]. For example, the corresponding theorem for `z_EVF` function is as follows:

Theorem 2.

```
⊢ ∀ E L m n P.
        ∼(NULL E) ∧ ∼(NULL L) ∧ m < LENGTH L ∧ n < LENGTH L ∧
        P < LENGTH E ∧ (∀ a. MEM a E ⇒ ∼MEM (FST a) [HD L])
        ⇒ z_EVF L E P m n = (F,TL L)
```

The assumptions of above theorem ensure that both lists `L` and `E` are not empty and the arguments of the function `z_EVF` are bounded by the `LENGTH` of `L` and `E`. The last conjunct in the assumption list models the no-reaction-possibility condition in the context of the function `z_EVF`. The conclusion of the theorem states that no update takes place under the given conditions by ensuring that the function `z_EVF` returns a pair with the first element being `F` (False), representing no match, and the second element being equal to `TL L`, which is actually equal to the original list `L` since an element was appended on head of `L` by the parent function.

4.2 Scenario 2: Single Reaction

The second scenario complements the first scenario and caters for the case when a reaction is possible and we verify that the molecules list is indeed updated based on the outcomes of that reaction. In order to be able to track the reaction and the corresponding update, we limit ourselves to only one reaction in this scenario but since we verify a generic theorem (universally quantified) for all possibilities our result can be extended to cater for multiple reactions as well. The theorem corresponding to this scenario 2 is as follows:

Theorem 3.

```
⊢ ∀ E L z m' n'.
      ∼NULL E ∧ ∼NULL (SND (EL z E)) ∧ 1 < LENGTH L ∧
      m' ≠ n' ∧ m' < LENGTH L ∧ n' < LENGTH L ∧ z < LENGTH E ∧
      ALL_DISTINCT (L ++ SND (EL z E)) ∧
      (∀ a b. a ≠ b ⇒ FST (EL a E) ≠ FST (EL b E)) ∧
      (∀ K m n. m < LENGTH K ∧ n < LENGTH K ∧
      (∀ j. MEM j K ⇒ MEM j L ∨ ∃ q. MEM q E ∧ MEM j (SND q)) ⇒
        if (EL m K = EL m' L) ∧ (EL n K = EL n' L)
          then HD (z_conj_int K m n) = FST (EL z E)
          else ∀ a. MEM a E ⇒ FST a ≠ HD (z_conj_int K m n))
      ⇒ z_deduction L E = z_del (L ++ SND (EL z E)) m' n'
```

The first two assumptions ensure that neither the list `E`, i.e., the list of EVFs, nor the second element of the pair at index z of the list `E` is empty. Similarly, the third assumption ensures that the list `L`, i.e., the list of initial molecules, contains at least two elements. These constraints ensure that we can have at least one reaction with the resultant being available at index `z` of the EVF list. The next four assumptions ensure that the indices `m'` and `n'` are distinct and these along with the index `z` fall within the range of elements of their respective lists of molecules `L` or EVFs `E`. According to the next assumption, i.e., `ALL DISTINCT`

(L ++ SND (EL z E)), all elements of the list L and the resulting molecules of the EVF at index z are distinct, i.e., no molecule can be found two or more times in the initial list L or the post-reaction list E. The next assumption, i.e., (∀ a b. a ≠ b ⇒ FST (EL a E) ≠ FST (EL b E)), guarantees that all first elements of the pairs in list E are also distinct. Note that this is different from the previous condition since the list E contains pairs as elements and the uniqueness of the pairs does not ensure the uniqueness of its first elements. The final condition models the presence of only one pair of reactants scenario. According to the assumptions of this implication condition, the variable K is used to represent a list that only has elements from list L or the second elements of the pairs in list E. Thus, it models the molecules list in a live experiment. Moreover, the variables m and n represent the indices of the list K and thus they must have a value less than the total elements in the list K (since the first element is indexed 0 in the HOL4 formalization of lists). Now, if the indices m and n become equal to m' and n', respectively, then the head element of the z_conj_int K m n would be equal to FST of EL z E. Otherwise, for all other values of indices m and n, no combination of molecules obtained by HD(Z_conj_int K m n) would be equal to the first element of any pair of the list E. Thus, the if case ensures that the variables m' and n' point to the reacting molecules in the list of molecules L and the variable z points to their corresponding resultant molecule in the EVF list. Moreover, the else case ensures that there is only one set of reacting molecules in the list L. The conclusion of the theorem formally describes the scenario when the resulting element, available at the location z of the EVF list, is appended to the list of molecules while the elements available at the indices m' and n' of L are removed during the execution of the function z_deduction on the given lists L and E.

The proof of Theorem 3 is again based on verifying sub-goals corresponding to this scenario for all the sub-functions, i.e., z_EVF, z_recur1, z_recur2 and z_deduction_recur. The formal reasoning for all of these proofs involved various properties of the del function for a list element and some of the key theorems developed for this purpose in our development are given in Table 3 and more details can be found in [1].

The formalization described in this section consumed about 500 man hours and approximately 2000 lines of HOL4 code, mainly due to the undecidable nature of higher-order logic. However, this effort raises the confidence level on the correctness of our formalization of Zsyntax . This fact distinguishes our work from all the other formal methods based techniques used in the context of BRNs, where the deduction rules are applied without being formally checked. Moreover, our formally verified theorems can also be used in the formal analysis of molecular pathways. The assumptions of these theorems provide very useful insights about the constraints under which a reaction or no reaction would take place. To the best of our knowledge, this is the first time that properties, like Theorems 1 and 3, about a molecular pathway experiment have been formally verified. Thus, the identification of these properties and their formal verification both constitute contributions of this paper.

Table 3. Formally Verified Properties of the `del` Function

Signature	Theorem
del_ASSOC_THM	$\vdash \forall$ L E m. m $<$ LENGTH L $\Rightarrow$ del (L ++ E) m = del L m ++ E
del_LENGTH_THM	$\vdash \forall$ L E m. m $<$ LENGTH L $\Rightarrow$ LENGTH (del L m) = LENGTH L $-$ 1
del_EL_THM	$\vdash \forall$ L m n. m $<$ n $\wedge$ n $<$ LENGTH L $\wedge$ 1 $<$ LENGTH L $\Rightarrow$ EL m L = EL m (del L n)
del_DISTINCT_THM	$\vdash \forall$ L n. n $<$ LENGTH L $\wedge$ ALL_DISTINCT L $\Rightarrow$ ALL_DISTINCT (del L n)
del_MEM_THM	$\vdash \forall$ L a m. m $<$ LENGTH L $\wedge$ MEM a (del L m) $\Rightarrow$ MEM a L
del_NOT_MEM_THM	$\vdash \forall$ L m. ALL_DISTINCT L $\wedge$ m $<$ LENGTH L $\Rightarrow\sim$ MEM (EL m L) (del L m)

5 Case Study: Pathway leading to Fructose-1,6-bisphosphate

Formation of Fructose-1,6-bisphosphate (F1,6P) is an intermediate step in glycolysis, i.e., a sequence of enzyme catalyzed reaction that breaks down glucose and forms pyruvate, which is then used to supply energy to living cells through the citric acid cycle [18]. In this section, we show how this pathway involving F1,6P can be formally verified in HOL4 using our formalization of Zsyntax .

The theorem representing the reaction of the glycolytic pathway leading from D-Glucose to F1,6P [4] can be described in classical Zsyntax format as follows:

```
Glc & HK & GPI & PFK & ATP & ATP ⊢ F1,6P
```

Using our formalization, this theorem can be defined in HOL4 as follows:

```
⊢ DISTINCT [Glc; HK; GPI; PFK; ATP; ADP; G6P; F6P; F16P] ⟹
(z_conj_elim (z_deduction [[Glc];[HK];[GPI];[PFK];[ATP];[ATP]]
                          [([Glc;HK],[[HK;Glc]]);
                           ([HK;Glc;ATP],[[HK];[G6P];[ADP]]);
                           ([G6P;GPI],[[F6P];[GPI]]);
                           ([F6P;PFK],[[PFK;F6P]]);
                           ([PFK;F6P;ATP],[[PFK];[F16P];[ADP]])] ) [F16P]
= [[F16P]]
```

The first list argument of the function `z_deduction` is the initial aggregate (IA) of molecules that are available for reaction and the second list argument of the function `z_deduction` represents the valid EVFs for this reaction. The EVFs mentioned in the form of pairs and involving the molecules (G6P, F6P, etc.) are obtained from wet lab experiments, as reported in [4]. The `DISTINCT` function used above makes sure that all molecule variables (from initial aggregate and EVFs) used in this theorem represent distinct molecules. Thus, the function

`z_deduction` would deduce the final list of molecules under these particular conditions. The function `z_conj_elim` will return the molecule `F1,6P` if it is present in the post-reaction list of molecules, as previously described.

Figure 1 shows the pathway leading to F1,6P in a step-wise manner. The gray-coloured circles show the chemical interactions and black colour represents the desired product in the pathway, whereas each rectangle shows total number of molecules in the reaction at a given time. It is obvious from the figure that whenever a reaction yields a product, the reactants get consumed (no longer remain in the list) hence satisfying the stoichiometry of a reaction.

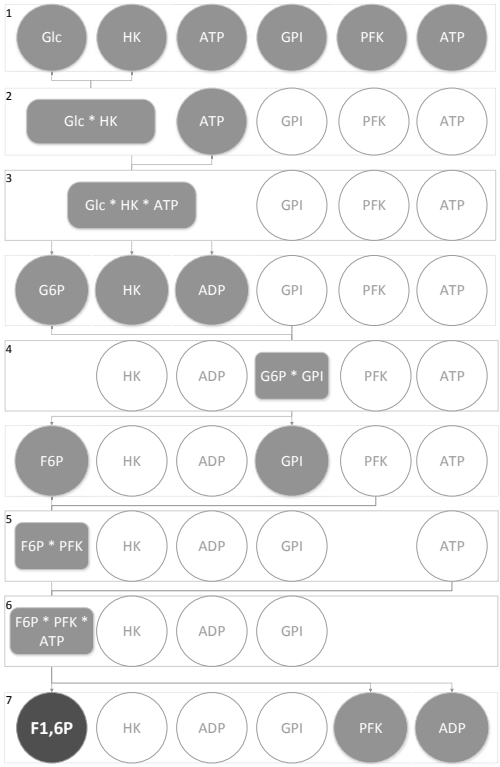

Fig. 1. Reaction Representing the Formulation of F1,6P

As part of this work, we also developed a simplifier `Z_SYNTAX_SIMP` [1] that simplifies the proof with a single iteration of the function `z_deduction_recur` and works very efficiently with the proofs involving our functions. The proof steps can be completely automated and the proof can be done in one step as well. However, we have kept the reasoning process manual purposefully as this way users can observe the status of the reaction at every iteration, which is a

very useful feature to get an insight of what is happening inside a reaction. Each application of `Z_SYNTAX_SIMP` on the reaction, depicted in Figure 1, would result in moving from a state n to $n + 1$.

The verification time required for each iteration step is given in Table 4. HOL4 was running on a linux based machine (Intel Core i5, 4GB RAM). The iteration time depends on the total number of molecules (elements of list) present at a given iteration. Low number of molecules translate to less number of possible combinations, which in turn leads to less time required to move to the next iteration.

Table 4. Runtime per Iteration

Iteration	**Duration (Seconds)**
$1 \rightarrow 2$	11.996
$2 \rightarrow 3$	7.376
$3 \rightarrow 4$	12.964
$4 \rightarrow 5$	12.756
$5 \rightarrow 6$	9.240
$6 \rightarrow 7$	0.048

Our HOL4 proof script is available for download [1], and thus can be used for further developments and analysis of different molecular pathways. It is important to note that formalizing Zsyntax and then verifying its properties was a very tedious effort. However, it took only 10 lines of code to define and verify the theorem related to the above case study in HOL4, which clearly illustrates the usefulness of our foundational work.

We have shown that our formalization is capable of modeling molecular reactions using Zsyntax inference rules, i.e., given a set of possible EVFs, our formalism can derive a final aggregate **B** from an initial aggregate **A** automatically. In case of a failure to deduce **B**, the proposed method still provides the biologist with all the intermediate steps so that one can examine the reaction in detail and figure out the possible cause of failure.

The evident benefit of our reasoning approach is its automatic nature as the user does not need to think about the proof steps and which EVFs to apply where. However, the most useful benefit of the proposed approach is its accuracy as the theorems are being verified in a formal way using a sound theorem prover. Thus, there is no risk of human error or wrong application of EVFs. Finally, due to the computer-based analysis, the proposed approach is much more scalable than the paper-and-pencil based analysis presented in [4].

6 Conclusion

Most of the existing formal verification research related to molecular biology has been focussed on using model checking. As a complementary approach, the

primary focus of the current paper is on using a theorem prover for reasoning about molecular pathways. The main strength of this approach, compared to existing model checking related work, is that the underlying methods and deduction rules can also be formally verified besides the verification of a particular molecular pathway case. Leveraging upon this strength, we formally verified two key behavioral properties of molecular pathways based on the Zsyntax language, which presents a deduction style formalism for molecular biology in the most biologist-centered way. Besides ensuring the correctness of our formalization of the Zsyntax operators and inference rules, the formally verified properties also play a vital role in reasoning about molecular pathways in the sound core of a theorem prover. The practical utilization and effectiveness of the proposed development has been shown by presenting the automatic analysis of Glycolytic pathway leading to Fructose-1,6-bisphosphate.

The proposed work opens the doors to many new directions of research. Firstly, we are developing a GUI to add more biologist friendly features in it. Moreover, we are also targeting some larger case studies, such as Dysregulation of the cell cycle pathway during tumor progression [16] and Fanconi Anemia/Breast Cancer (FA/BRCA) pathway [26]. Another interesting future direction is to leverage the high expressiveness of higher-order-logic and utilize calculus and differential theoretic reasoning to add reaction kinetics support in our formalism.

References

1. S. Ahmad. Formal Reasoning about Molecular Pathways - HOL Proof Script. `http://save.seecs.nust.edu.pk/projects/holsyntax/holzsyntax.html`, 2014.
2. S. Ahmad, O. Hasan, and U. Siddique. Towards Formal Reasoning about Molecular Pathways in HOL. In *International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises*, pages 378–383. IEEE, 2014.
3. C. Baier and J. Katoen. *Principles of Model Checking*. MIT Press, 2008.
4. G. Boniolo, M. D'Agostino, and P. Di Fiore. Zsyntax: a Formal Language for Molecular Biology with Projected Applications in Text Mining and Biological Prediction. *PloS ONE*, 5(3):e9511–1–e9511–12, 2010.
5. A. Church. A Formulation of the Simple Theory of Types. *Journal of Symbolic Logic*, 5(2):56–68, 1940.
6. V. Danos and C. Laneve. Formal Molecular Biology. *Theoretical Computer Science*, 325(1):69–110, 2004.
7. N.H. Hunt et al. Immunopathogenesis of Cerebral Malaria. *International Journal for Parasitology*, 36(5):569–582, 2006.
8. L. Trilling Fab. Corblin, E. Fanchon. Applications of a Formal Approach to Decipher Discrete Genetic Networks. *BMC Bioinformatics*, 11(1):385, 2010.
9. W. Fontana. Systems Biology, Models, and Concurrency. *SIGPLAN Notices*, 43(1):1–2, January 2008.
10. F. Cassez et al. G.Bernot. Semantics of Biological Regulatory Networks. *Electronic Notes Theoretical Computer Science*, 180(3):3–14, 2007.
11. Peter J. E. Goss and J. Peccoud. Quantitative Modeling of Stochastic Systems in Molecular Biology by using Stochastic Petri Nets. *Proceedings of the National Academy of Sciences*, 95(12):6750–6755, 1998.

12. J. Harrison. *Handbook of Practical Logic and Automated Reasoning*. Cambridge University Press, 2009.
13. K. Hirayama. Genetic Factors Associated with Development of Cerebral Malaria and Fibrotic Schistosomiasis. *Korean J. Parasitol*, 40(4):165–172, Dec 2002.
14. M. Magnin L. Paulevé and O. Roux. Abstract Interpretation of Dynamics of Biological Regulatory Networks. *Electronic Notes Theoretical Computer Science*, 272(0):43–56, 2011.
15. C.J. Langmead. Generalized Queries and Bayesian Statistical Model Checking in Dynamic Bayesian Networks: Application to Personalized Medicine. In *Proc. International Conference on Computational Systems Bioinformatics*, pages 201–212, 2009.
16. R. Maglietta, V. Liuzzi, E. Cattaneo, E. Laczko, A. Piepoli, A. Panza, M. Carella, O. Palumbo, T. Staiano, F. Buffoli, A. Andriulli, G. Marra, and N. Ancona. Molecular Pathways Undergoing Dramatic Transcriptomic Changes During Tumor Development in the Human Colon. *BMC Cancer*, 12(1):608, 2012.
17. R. Milner. A Theory of Type Polymorphism in Programming. *Journal of Computer and System Sciences*, 17(3):348–375, 1977.
18. D. Nelson. *Lehninger Principles of Biochemistry*. W.H. Freeman, New York, 2008.
19. L.C. Paulson. *ML for the Working Programmer*. Cambridge University Press, 1996.
20. R. Pelánek. Fighting State Space Explosion: Review and Evaluation. In *Formal Methods for Industrial Critical Systems*, volume 5596 of *Lecture Notes in Computer Science*, pages 37–52. Springer, 2008.
21. J. Pospchal and V. Kvasnika. Reaction Graphs and a Construction of Reaction Networks. *Theoretica Chimica Acta*, 76(6):423–435, 1990.
22. A. Regev and E. Shapiro. Cells as Computation. *Nature*, 419:343, 2002.
23. A. Regev and E. Shapiro. The π-Calculus as an Abstraction for Biomolecular Systems. In *Modelling in Molecular Biology*, Natural Computing Series, pages 219–266. Springer, 2004.
24. A. Regev, W. Silverman, and E. Y. Shapiro. Representation and Simulation of Biochemical Processes Using the pi-Calculus Process Algebra. In *Pacific Symposium on Biocomputing*, pages 459–470, 2001.
25. M. Rizzotti and A. Zanardo. Axiomatization of Genetics. 1. Biological Meaning. *Journal of Theoretical Biology*, 118(1):61–71, 1986.
26. A. Rodríguez, D. Sosa, L. Torres, B. Molina, S. Frías, and L. Mendoza. A Boolean Network Model of the FA/BRCA Pathway. *Bioinformatics*, 28(6):858–866, 2012.
27. L. Thomas and R. d' Ari. *Biological Feedback*. CRC Press, USA, 1990.
28. R. Thomas. *Kinetic Logic: A Boolean Approach to the Analysis of Complex Regulatory Systems*, volume 29 Lecture Notes in Biomathematics. Springer-Verlag, 1979.
29. J.J Tyson, C.A. Nagy, and B. Novak. The Dynamics of Cell Cycle Regulation. *Bioessays*, 24(12):1095–1109, 2002.
30. O. Wolkenhauer, D. Shibata, and M.D. Mesarovic. The Role of Theorem Proving in Systems Biology. *Journal of Theoretical Biology*, 300(0):57–61, 2012.
31. J.H. Woodger, A. Tarski, and W.F. Floyd. *The Axiomatic Method in Biology*. The University Press, 1937.
32. A. Zanardo and M. Rizzotti. Axiomatization of Genetics 2. Formal Development. *Journal of Theoretical Biology*, 118(2):145–152, 1986.

Towards a Family of Test Selection Criteria for Symbolic Models of Real-Time Systems

Diego R. Almeida[1], Alan Moraes[2,3],
Wilkerson L. Andrade[2], and Patrícia D. L. Machado[2]

[1] IFPE, Afogados da Ingazeira, PE, Brazil
diego.rodrigues@afogados.ifpe.edu.br
[2] Software Practices Laboratory (SPLab), UFCG, Campina Grande, PB, Brazil
{wilkerson,patricia}@computacao.ufcg.edu.br
[3] Informatics Center, UFPB, João Pessoa, PB, Brazil
alan@ci.ufpb.br

Abstract. In model-based testing, test cases are generated from a specification model. To avoid an exhaustive search for all possible test cases that can be obtained, usually an expensive and infeasible activity, test case generation may be guided by a test selection criterion. The objective of a test selection criterion is to produce a minimal test suite and yet effective to reveal faults. However, the choice of a criterion is not straightforward specially for real-time systems, because most criteria presented in the literature are general-purpose. Moreover, the relationship between general-purpose and specific criteria for real-time systems is not clear. In this paper, we investigate the criteria that can be applied for test case generation in the scope of model-based testing of real-time systems, specifically of Timed Input-Output Symbolic Transition Systems (TIOSTS) models. We formalize a family of 19 test selection criteria ordered by strict inclusion relation for TIOSTS models. The family combines general-purpose data-flow-oriented and transition-based criteria with specific reactive and real-time systems criteria. We also perform an empirical study to compare the effectiveness of selected criteria. Results of the empirical study indicate that failure detection capability of the generated test suite may vary, but differences are not significant for time failures. We conclude that more effective criteria for the model-based testing of real-time systems are still needed.

1 Introduction

Model-Based Testing is a testing approach that relies on the design of abstract models of an application to generate, execute and evaluate tests [10, 22, 27]. It has been applied with success in industry, with special emphasis in the avionic, railway and automotive domains [21].

Test case generation algorithms are based on test selection criteria that guide how to search for test cases and when to stop the test case generation process. Different test suites can be generated depending on the chosen test selection criterion [29]. They may vary in size, behavior coverage and failure detection

capability. While it is more likely that larger (and possibly with higher model coverage) test suites have better failure detection capability than smaller (and possibly with lower model coverage) ones, they are usually more expensive to manage and to execute. Therefore, test selection criteria need to establish how to guarantee the generation of test suites that are ultimately cost-effective.

Real-time systems are reactive systems whose behavior is constrained by time [18]. They usually combine concurrent execution of processes, consequently the nature of their failures is complex. The testing of these systems should uncover time-related faults that may require specific test cases to be exercised.

Most test selection criteria for real-time systems at model level are based on structural elements of a model behavior and its data usage [14]. Some specific test selection criteria for real-time systems have been proposed, such as covering all clock resets and all guard bounds [12]. However, the choice of a criterion is not straightforward, because the relationship between general-purpose and specific criteria for real-time systems is not clear [2].

In this paper, we investigate test selection criteria for real-time systems in the context of model-based testing. We focus on criteria that can be applied to transition systems, because they are usually the basis for conformance testing of real-time systems [17,28]. We use Timed Input-Output Symbolic Transition Systems (TIOSTS) models [5,6], where system behavior is modeled as a transition system with data and time symbolically defined.

This paper makes two contributions. First, we formalize a family of 19 criteria partially ordered by strict inclusion relation for TIOSTS models. The family combines Transition-Based Criteria, data-flow-oriented criteria, Reactive Systems Criteria and Real-Time Systems Criteria. We prove inclusion or incompatibility whenever our family diverges from the known relationship in other models, because some relation between criteria change when applied to TIOSTS models.

Second, we conduct a controlled experiment to compare the effectiveness of selected criteria. The empirical study measures the size, the failure detection capability and the rate of failures detected by the size of the test suite of different criteria. In order to conduct the empirical study, we implemented a selection of criteria from the family using a depth-first search-based algorithm. Statistical analyses show that the criteria present different failure detection capability, although, significant differences cannot be observed for time-related failures. Furthermore, current specific criteria for real-time systems lack precision, i.e. they miss important failures, pointing to the need for further research in this area.

The paper is structured as follows. Section 2 introduces the TIOSTS model and test selection criteria for model-based testing of real-time systems. Section 3 formalizes a family criteria for TIOSTS. Section 4 presents an empirical study to compare selected criteria. Section 5 discusses related work. Finally, Section 6 presents concluding remarks along with pointers for further research.

2 Background

This section presents the symbolic model on which this work is based and introduces the concept of test selection criterion in the context of model-based testing.

2.1 Timed Input-Output Symbolic Transition System Model

Timed Input-Output Symbolic Transition System (TIOSTS) [5, 6] is a symbolic model for real-time systems that handles both data and time. The TIOSTS model was defined as an extension of two existing models: Timed Automata [3] and Input-Output Symbolic Transition Systems [15, 24]. Basically, a TIOSTS is an automaton with a finite set of locations where system data and time evolution are respectively represented by variables and a finite set of clocks. The transitions of the model are composed of a guard on variables and clocks, an action with parameters, an assignment of variables, and a set of clocks to reset.

Figure 1 shows an example of TIOSTS that models a machine for refilling a card for using the subway. Initially, the system is in the `Idle` location where it expects the `Credit` input carrying the desired `value` to refill, then this value is saved into the `refillValue` variable[4] and `balance` is initialized to zero.

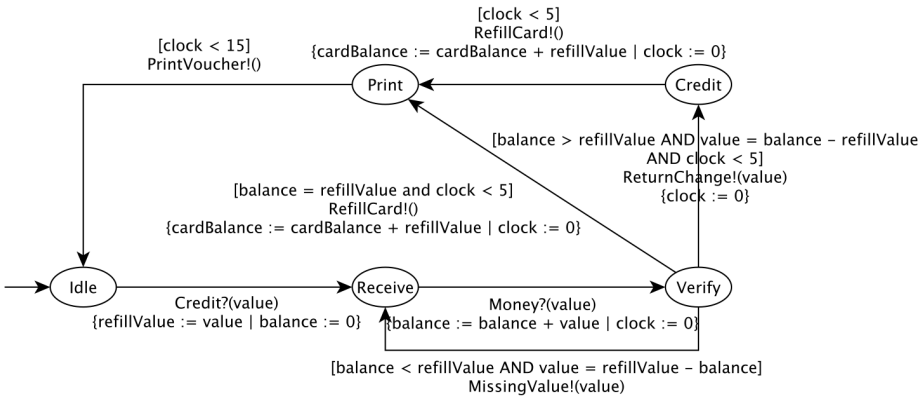

Fig. 1. TIOSTS model of a refilling machine

From the `Receive` location to `Verify` the client informs the amount to be credited to the card. This value is accumulated in the `balance` variable and the `clock` is set to zero. If the current balance is less than the desired value to refill, then the `Receive` location is reached again and the `MissingValue` output is emitted for informing the remaining value (the condition `value` =

[4] Action parameters have local scope, thus their values must be stored in variables for future references.

`refillValue` − `balance` contained in the guard means "choose a value for the `value` parameter that, with the values of `refillValue` and `balance` variables, satisfies the guard").

From the `Verify` location, if the `balance` is greater than `refillValue` some value must be returned to the client in less than 5 time units. After that, the `clock` is reset to zero again. Then, the `RefillCard` output action must be performed in less than 5 time units and the `cardBalance` is increased by `refillValue`. Otherwise, from `Verify`, if `balance` is exactly equals to `refillValue` the card must be refilled in less than 5 time units. Finally, from the `Print` location, the voucher must be printed in less than 15 time units and `Idle` location is reached again. A formal definition of TIOSTS models is presented in Definition 1 [5].

Definition 1 (TIOSTS). *A TIOSTS is a tuple* $W = \langle V, P, \Theta, L, l^0, \Sigma, C, \mathcal{T} \rangle$, *where:*

- V *is a finite set of typed variables;*
- P *is a finite set of parameters. For* $x \in V \cup P$, $type(x)$ *denotes the type of* x*;*
- Θ *is the initial condition, a predicate with variables in* V*;*
- L *is a finite, non-empty set of locations and* $l^0 \in L$ *is the initial location;*
- $\Sigma = \Sigma^? \cup \Sigma^!$ *is a non-empty, finite alphabet, which is the disjoint union of a set* $\Sigma^?$ *of input actions and a set* $\Sigma^!$ *of output actions. For each action* $a \in \Sigma$, *its signature* $sig(a) = \langle p_1, ..., p_n \rangle$ *is a tuple of distinct parameters, where each* $p_i \in P$ $(i = 1, ..., n)$*;*
- C *is a finite set of clocks with values in the set of non-negative real numbers, denoted by* $\mathbb{R}^{\geq 0}$*;*
- $\mathcal{T}$ *is a finite set of transitions. Each transition* $t \in \mathcal{T}$ *is a tuple* $\langle l, a, G, A, y, l' \rangle$, *where:*
  - $l \in L$ *is the origin location of the transition,*
  - $a \in \Sigma$ *is the action,*
  - $G = G^D \wedge G^C$ *is the guard, where* G^D *is a predicate over variables in* $V \cup set(sig(a))$[5,6] *and* G^C *is a clock constraint over* C *defined as a conjunction of constraints of the form* $\alpha \# c$, *where* $\alpha \in C$, $\# \in \{<, \leq, =, \geq, >\}$, *and* $c \in \mathbb{N}$,
  - $A = (A^D, A^C)$ *is the assignment of the transition. For each variable* $x \in V$ *there is exactly one assignment in* A^D, *of the form* $x := A^{D^x}$, *where* A^{D^x} *is an expression on* $V \cup set(sig(a))$. $A^C \subseteq C$ *is the set of clocks to be reset,*
  - $y \in \{lazy, delayable, eager\}$ *is the deadline of the transition,*
  - $l' \in L$ *is the destination location of the transition.* ◇

The semantics of a TIOSTS is described by Andrade and Machado [5]. Next we define the concepts of *state*, *path* and *test case*.

[5] G^D is assumed to be expressed in a theory in which satisfiability is decidable.
[6] Let $set(j)$ be the function that converts the tuple j in a set.

Definition 2 (State of TIOSTS). *In TIOSTS model, a state is a tuple* $\langle l, v_1, ..., v_n, c_1, ..., c_m \rangle$*, which consists of a location* $l \in L$*, a specific valuation for all variables* $v_i \in V$*, and a valuation for all clocks* $c_i \in C$*.* $\diamond$

Definition 3 (Path). *A path is a finite sequence of transitions* $(t_1, ..., t_k)$*,* $k \geq 1$*, such that the destination location of transition* t_i *is equal to the origin location of the transition* t_{i+1} *for* $i = 1, 2, ..., k-1$*.* $\diamond$

Definition 4 (Test Case). *A test case is a deterministic TIOSTS* $TC = \langle V_{TC}, P_{TC}, \Theta_{TC}, L_{TC}, l^0_{TC}, \Sigma_{TC}, C_{TC}, \mathcal{T}_{TC} \rangle$*, where* $\Sigma^?_{TC} = \Sigma^!_{\mathcal{S}}$ *and* $\Sigma^!_{TC} = \Sigma^?_{\mathcal{S}}$ *(actions are mirrored w.r.t. specification), equipped with three disjoint sets of verdict locations* Pass, Fail*, and* Inconclusive*. Furthermore, each sequence from the initial location* l^0_{TC} *to some verdict location is a path.* $\diamond$

According to Definition 4, the execution of a test case can emit one of three possible verdicts: *Pass*, *Fail*, and *Inconclusive*. *Pass* means that some targeted behavior of the system under test has been reached, *Fail* means rejection of the SUT, and *Inconclusive* means that targeted behavior cannot be reached anymore.

Figure 2 is a test case for the TIOSTS model of the refilling machine. The test case aims to exercise the scenario where the system emits the `RefillCard` output when the amount to be credited to the card (`value_2`) is equal to desired value to refill (`value_1`). In this case, the verdict is *Pass*. If the amount to be credited to the card (`value_2`) is less than the desired value to refill (`value_1`), and the system emits the `MissingValue` output with parameter equals to `value_1` − `value_2`, then the verdict is *Inconclusive*. It is *Inconclusive* because this behavior is specified in the model, but it is not the scenario the tester would like to observe in the test case execution. The same applies to `ReturnChange` output action of the test case. All other cases lead to the implicit *Fail* verdict.

2.2 Test Selection Criteria for Real-Time Systems

In model-based testing, test cases are derived from a model which specifies the expected behavior of a system under test. A *Test Selection Criterion* defines which parts of the system are going to be tested, how often and under what circumstances they will be tested [29]. Test selection criteria are used for two main purposes: to measure the adequacy of the test suite with respect to the level of quality required by the context, and to stop the test generation process after the criterion is reached [29].

We conducted a systematic literature review to identify studies that address test selection criteria for real-time systems at model level [2]. We considered studies that a criterion was used at least as part of a test case generation process in the scope of transition and state-based systems [1,7,9,12–14,16,17,20,26,31].

The results of the review show that most general-purpose test selection criteria may be applied to models of real-time systems. There are also specific criteria for real-time systems proposed in the literature. However, there is a lack

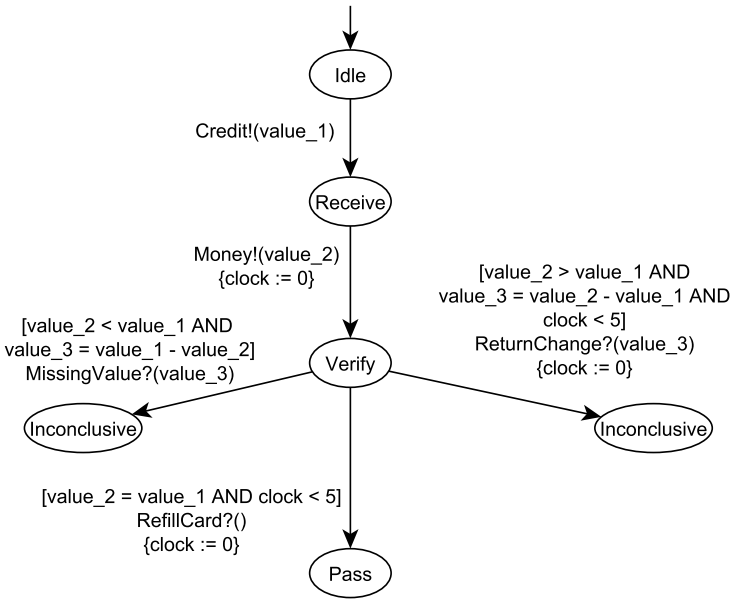

Fig. 2. A test case for the refilling machine.

of studies that investigate the theoretical and empirical relationship between criteria. The theoretical relationship could indicate the relative effort to satisfy a criterion, while the empirical evaluation could compare criteria effectiveness with respect to failure detection capability.

En-Nouaary [12] proposes a family of test selection criteria ordered by strict inclusion relation for Timed Input-Output Automata (TIOA). His family combines Transition-Based Criteria, Reactive Systems Criteria, and Real-Time Systems Criteria. But data-related criteria are not included because the TIOA model does not support data abstraction. Conversely, the TIOSTS model symbolically abstracts both time and data, thus data-related criteria can be applied to it. Furthermore, to the best of our knowledge, there is no work on test selection criteria for real-time systems at model level that evaluate the ability to reveal faults of selected criteria.

3 Towards a Family of Test Selection Criteria for TIOSTS

In this section, we propose a family of test selection criteria for TIOSTS models. We extend En-Nouaary's family [12] to include data-related criteria. We choose to include Data-Flow-Oriented Criteria, because they can be empirically evaluated with the same failure model employed to compare Transition-Based

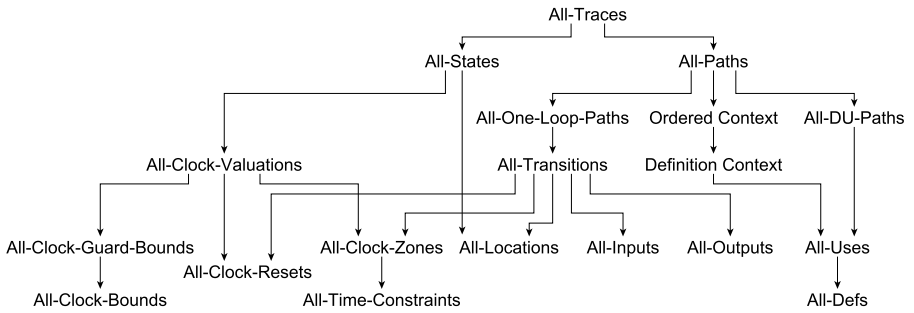

Fig. 3. Family of test selection criteria ordered by strict inclusion relation for TIOSTS models.

Criteria and Real-Time Systems Criteria in the next section. Thus our proposed family of criteria combines Transition-Based Criteria, Reactive Systems Criteria, Real-Time Systems Criteria and Data-Flow-Oriented Criteria. Table 1 describes the criteria we considered in this work.

Test selection criteria are often theoretically compared to each other by three relations: *strict inclusion*, *equivalence*, or *incompatibility* [23]. The relations are formalized in Definitions 6, 7 and 8 respectively.

Definition 5 (Inclusion Relation). *A criterion* c_1 includes *a criterion* c_2 *if any set of test cases that satisfies* c_1 *also satisfies* c_2 *[23].* ◇

Definition 6 (Strict Inclusion Relation). *A criterion* c_1 strictly includes c_2*, denoted by* $c_1 \Rightarrow c_2$*, if* c_1 *includes* c_2 *but there is a set of test cases that satisfies* c_2 *but does not satisfy* c_1*. Note that this is a transitive relation [23].* ◇

Definition 7 (Equivalence Relation). *A criterion* c_1 *is* equivalent *to a criterion* c_2 *if* c_1 *includes* c_2 *and* c_2 *includes* c_1*.* ◇

Definition 8 (Incompatible Relation). *A criterion* c_1 *is* incompatible *with a criterion* c_2 *if* c_1 *does not include* c_2 *and* c_2 *does not include* c_1*.* ◇

Our goal is to produce a sound family of test selection criteria partially ordered by strict inclusion relation. We do not intend to prove all equivalences or incompatibilities between criteria. To accomplish this, our strategy is i) to reuse the proofs of strict inclusion relations from other formalisms if they are also valid for TIOSTS; ii) to prove new strict inclusion relations resulting from the combination of classes of criteria; iii) to prove the exclusion of strict inclusion relations valid for other formalisms but not valid for TIOSTS. The proposed family is formalized in Theorem 1.

Theorem 1. *The family of criteria for TIOSTS is partially ordered by strict inclusion as shown in Figure 3. Furthermore,* $c_1 \Rightarrow c_2$ *iff it is explicitly shown to be so in Figure 3 or follows from the transitivity of the relationship.*

Table 1. Test Selection Criteria for TIOSTS models.

Criterion	Description
Transition-Based Criteria	
All-Locations [12, 16]	Every location of the model must be exercised by at least one test case.
All-Paths [12, 14]	Every path of the model must be exercised by at least one test case.
All-One-Loop-Paths [29]	Every loop-free paths through the model must be exercised, plus all the paths that loop at least once.
All-Transitions [12, 16]	Every transition of the model must be exercised by at least one test case.
All-States [1, 12, 31]	Every state of the model must be exercised by at least one test case.
All-Traces [12]	Every trace of the model must be included in the test suite.
Data-Flow-Oriented Criteria	
All-Defs [29]	At least one def-use pair(d_v, u_v) for every definition d_v must be exercised by at least one test case, i.e. at least one path from every definition to one of its use must be covered.
All-DU-Paths [29]	Every path for all def-use pairs(d_v, u_v) must be exercised by at least one test case, i.e. all paths from every definition d_v to every use u_v must be covered.
All-Uses [29]	Every def-use pairs(d_v, u_v) must be exercised by at least one test case, i.e. at least one path from every definition d_v to every use u_v must be covered.
Definition Context [14]	All paths from every context of definition of variable x to the definition of variable x must be exercised by at least one test case. The context of definition of the variable x are the transitions where the variables used to define the value of x are defined.
Ordered Context [14]	Similar to Definition Context, but the transitions context are listed in the order of their definitions.
Reactive Systems Criteria	
All-Inputs [9, 12]	Every input action of the model must be exercised by at least one test case.
All-Outputs [9, 12]	Every output action of the model must be exercised by at least one test case.
Real-Time Systems Criteria	
All-Clock-Bounds [12]	Every clock bound of the model must be exercised by at least one test case. The bound of a clock is the highest value that a clock can assume.
All-Clock-Guard-Bounds [12]	Every clock guard bound of the model must be exercised by at least one test case. This criterion is similar to All-Clock-Bounds but considering only the time guards.
All-Clock-Valuations [12]	Every clock valuation of the model must be exercised by at least one test case.
All-Clock-Resets [12]	Every clock reset of the model must be exercised by at least one test.
All-Clock-Zones [12, 26]	Every clock zone of the model must be visited through at least one test case, i.e. all transitions with clock resets or time guards must be covered.
All-Time-Constraints [12]	Every time guard of the model must be exercised by at least one test case.

Note: The criteria in this table are defined in terms of satisfiable paths, i.e. all data and time guards in a path must be satisfiable.

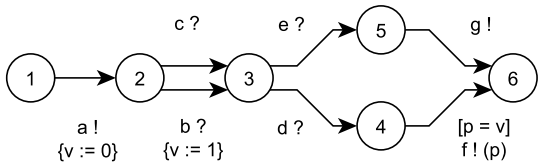

Fig. 4. A TIOSTS model to assist in the proof of All-DU-Paths $\not\Rightarrow$ All-Transitions.

Proof. We need to prove the relations All-States $\Rightarrow$ All-Locations, All-One-Loop-Paths $\Rightarrow$ All-Transitions, All-Transitions $\Rightarrow$ All-Clock-Resets, and All-DU-Paths $\not\Rightarrow$ All-Transitions. All other relations can be easily checked based on proofs already presented in the literature [12, 23, 29, 32].

1. All-States $\Rightarrow$ All-Locations. Proof follows directly from the definitions of the criteria. We recap that a state of a TIOSTS consists of a location, a specific valuation for all variables, and a valuation for all clocks. Since the All-States criterion demands the all states to be covered, thus All-States $\Rightarrow$ All-Locations.
2. All-One-Loop-Paths $\Rightarrow$ All-Transitions. Proof follows directly from the definitions of the criteria. The All-One-Loop-Paths criterion demands that all loop-free paths to be covered plus all loops at least one lap. Since all transitions must be either in a loop-free path or in a loop, thus All-One-Loop-Paths $\Rightarrow$ All-Transitions.
3. All-Transitions $\Rightarrow$ All-Clock-Resets. Proof follows directly from the definitions of the criteria. A clock reset happens within the assignment of a transition. The All-Transitions criterion demand that all transitions to be covered. Since all transitions with clock resets are a subset of all transitions, thus All-Transitions $\Rightarrow$ All-Clock-Resets.
4. All-DU-Paths $\not\Rightarrow$ All-Transitions. Proof by contradiction. Let's assume that All-DU-Paths $\Rightarrow$ All-Transitions. Consider the TIOSTS model in the Figure 4. The model has two def-use pairs: $\{(q_1, [true], a!, \{v := 0\}, \emptyset, q_2), (q_4, [p = v], f!(p), \emptyset, \emptyset, q_6)\}$ and $\{(q_2, [true], b?, \{v := 1\}, \emptyset, q_3), (q_4, [p = v], f!(p), \emptyset, \emptyset, q_6)\}$. The test cases[7] $\{\{a! \rightarrow c? \rightarrow d? \rightarrow f!(p)\}, \{a! \rightarrow b? \rightarrow d? \rightarrow f!(p)\}\}$ satisfy the All-DU-Paths criterion for this model, but the transitions $(q_3, true, e?, \emptyset, \emptyset, q_5)$ and $(q_5, true, g!, \emptyset, \emptyset, q_6)$ are not covered. Thus our assumption is incorrect, and All-DU-Paths $\not\Rightarrow$ All-Transitions. □

It is important to remark that the relation All-Uses $\Rightarrow$ All-Transitions does not hold for TIOSTS as it does for other models [23]. In fact, even All-DU-Paths $\not\Rightarrow$ All-Transitions for TIOSTS. This happens because a transition in TIOSTS may have neither a definition nor a use of a variable. Thus not all transitions will be covered by the All-DU-Paths criterion.

[7] The last transition in the test case leads to the *Accept* location.

4 Empirical Study

In this section we present a controlled experiment to compare the effectiveness of selected criteria. We follow the guidelines given by Wohlin, Runeson, Höst and Ohlsson [30]. The main goal of the empirical study is to investigate test selection criteria for real-time systems by observing the test suite generated from TIOSTS models according to a given criterion with respect to their size and failure detection capability from the point of view of the tester in the context of model-based testing. The research hypothesis is that different criteria may generate different suites of different sizes that may reveal a number of different failures.

Planning. We conducted this experiment in a research laboratory — an offline study with a specific context. As independent variable, we have the test selection criterion. The treatments are: All-One-Loop-Paths (AOLP), All-Transitions (AT), All-Locations (AL), All-Clock-Zones (ACZ), All-Clock-Resets (ACR), All-DU-Paths (ADUP), All-Uses (AU), and All-Defs (AD). Instead of evaluating all criteria of the family, we choose to evaluate the most used criteria found in our literature review. The selected criteria are representative of transition, time and data-related criteria.

The dependent variables are: i) size of the generated test suites ($Size$); and ii) failure detection capability, measured as the number of different failures that can be detected ($Failure$). From these dependent variables, for each treatment and object, we computed two values: i) the *percentage of failure*, defined as the relation between the $Failure$ value and the total of possible failures; ii) the *density of failure* as the relation between the $Failure$ and the $Size$ values. For the sake of simplicity, the hypotheses of the study are formulated based on these measures only as follows. Let $\%failure_i = \frac{Failure_i}{TotalFailures}$ and $density_i = \frac{Failure_i}{Size_i}$, where i is a test criterion and $Failure_i$, $Size_i$ are the average value of the correspondent dependent variables for each of the considered objects. Based on statistical testing, the null hypothesis is defined as the equality of all criteria, whereas the alternative hypothesis is defined as the difference between all criteria.

Regarding experimental design, this study consists of one factor and eight levels (eight test criteria) with six repetitions corresponding to six different models from three applications of real-time systems presented in the literature. We considered a confidence of 95% when deciding on hypothesis rejection. As input, for each criterion, only TIOSTS models are required. Dependent variables are computed automatically. Therefore, there is no human intervention and no subjects to be considered. Since there are no random choices involved, there is no need to compute the number of replications required.

The objects (TIOSTS models) were obtained from 3 different applications: i) Alarm System — Monitoring and actuation system that can detect invasion and also the presence of intruders in a building through door, window and movement sensors [25]; ii) Aircraft Attack System — System that controls attacks to specific land targets and also threat detection from a missile or another aircraft [19]; and

iii) Philips Audio Protocol — Protocol that defines control message exchanging for audio and video devices [8]. Moreover, collisions detection and delivery failure are handled. From these applications, we created six models and used them as input to the test case generator we implemented using a depth-first search-based algorithm. Table 2 presents the metrics of number of locations, transitions, transitions with time constraints, and transitions with data constraints of the considered models.

Table 2. Metrics of real-time system models used in the empirical study.

Model	Locations	Transitions	Trans. w/ time constraints	Trans. w/ data constraints
Alarm1	7	9	6	7
Alarm2	10	23	13	19
Aircraft1	11	13	8	6
Aircraft2	14	35	20	28
Protocol1	17	29	10	25
Protocol2	17	37	18	25

Notes. Alarm1: Alarm System without power failure. Alarm2: Simplified version of Alarm1 with power failure treatment. Aircraft1: Aircraft Attack System functionality only. Aircraft2: Simplified version of Aircraft1 with threat detection functionality. Protocol1: System without failure recovery. Protocol2: Simplified version of Protocol1 with failure recovery.

It is often difficult to associate a failure with a single fault at code level, because a failure may be caused by one or more faults. Therefore, for the purpose of this study and also to avoid undesired effects in the results, instead of the number of faults, we opt to measure failures — the number of different failures that can be detected by at least one test case in a given test suite. To allow for a reasonable sample of failures, we defined a failure model that contains potential failures which can be detected in a real-time system, particularly as a result of violation of time constraints. This model was based on previous studies such as the one performed by En-Nouaary, Khendek and Dssouli [11], and by Andrade and Machado [4]. Two basic types of failures were considered: time and behavior. The former is necessarily connected to non-conformity with time constraints, whereas the latter are more related to behavior non-conformity. For the sake of space, Table 3 presents only considered failures for the *Alarm2* model. Note that there is a different distribution of faults of the two types. The reason is that we do not aim to control this factor so that the distribution achieved is mostly a consequence of potential failures identified by considering each model.

Study execution was conducted according to the following process: 1) For each input model, a test suite was generated for each of the criteria; 2) For each test suite, each test case was analysed to determine whether it can fail according to the failure model; 3) For each test suite, failures from the failure model were marked when covered by the suite; 4) Data on study variables was collected; 5) *%failure* and *density* values were computed and analysis of results conducted.

Table 3. Failure Model for *Alarm2* model.

Failure	Type	Description
F04	Time	When power failure occurs, sensor status does not change.
F05	Time	When power failure is detected, the system does not change power supply on time.
F06	Behavior	After handling power failure, system does not resume execution as expected.
F07	Time	When power failure occurs, status change of movement sensor is not detected.
F08	Time	When power failure occurs, status change of window sensor is not detected.
F09	Behavior	After power failure handling, system does not detect an invaded room.
F10	Behavior	After power failure handling, alarm starts without invasion detection.

Threats to Validity. Measures were rigorously taken regarding data treatment and assumption with a confidence level of 95% that is usually applied in comparing studies. Also, to avoid the influence on the kind of applications in the obtained results, we have chosen specifications constructed by different authors — the models have different structural elements as illustrated in Table 2. Moreover, correctness of the implementation of the algorithms is critical to assess whether the results are reliable. Therefore, validation was throughly performed and, to avoid an inconsistent generation of suites, all algorithms are based on the same basic strategy — a depth-first search — where each criterion is applied as a stop condition. Furthermore, models used in the study may not be representative of all kinds of real-time systems, therefore, results can only be interpreted as specific. However, it is important to remark that they may be considered as an evidence since results confirm properties already known, particularly for the general criteria.

Results and Analysis. Data collected in the study as well as test cases generated can be downloaded from the study web site[8]. Figure 5 shows the box plots for the percentage of failure values and Figure 6 shows the box plot for the density of failures values. As the values do not follow a normal distribution, the Kruskal-Wallis test was performed and we obtained a *p-value* of 0.0388 for the percentage of failures. This means that we can reject the null hypotheses: when compared together the criteria present a different failure detection capability. However, if we consider only "Time" failures, the *p-value* would be 0.1487. Therefore, we can observe that, for the considered criteria, significant differences of capability for this kind of failure cannot be observed.

On the other hand, for the density of failure values, by applying the Kruskal-Wallis test we obtained a *p-value* of 0.0670. This means that we cannot reject the null hypotheses: we cannot observe a significant difference on the failure density for the considered criteria. It is also important to mention that no significant correlation between the values of size and failure has been observed for any of the considered criteria.

[8] https://sites.google.com/a/computacao.ufcg.edu.br/rtscoverage/

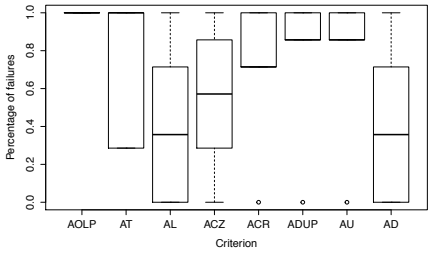

Fig. 5. Boxplot of percentage of failure detected for each criterion.

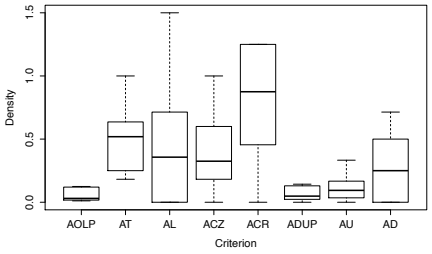

Fig. 6. Boxplot of density of failure detected for each criterion.

General Remarks. From this study, we can observe that the more general criteria such as All-One-Loop-Paths and All-Transitions as well as All-DU-Paths and All-Uses present a better failure coverage even when only time failures are considered. The reason is that more test cases are generated when these criteria are considered. However, they do not always present the best failure density capacity. Which means that a number of test cases may be either useless or redundant for the purpose of detecting the considered failures. From the general criteria, All-Uses (followed by All-DU-Paths) seems to present more consistently the best relation between size and failure detection capability. The reason is that they can most effectively explore the relation between events that are related to a given variable, whereas the structural criteria such as All-Transitions and All-Locations can miss certain combinations. The clock related criteria All-Clock-Zones and All-Clock-Resets present considerably smaller test suites and good density failure capacity, particularly the second one. However, not all failures are covered, even time related ones. Consequently, these criteria may only be considered under severe project constraints. Otherwise, one might consider using both of them together in order to improve failure detection capability and still keep a reasonable failure density.

5 Related Work

Test selection criteria for different kinds of models of real-time systems have already been investigated in the literature. But most of works just describe a criterion or a set of criteria without proper theoretical and empirical evaluation.

En-Nouaary [12] proposes a family of test selection criteria ordered by strict inclusion relation criteria for TIOA models. Our proposal is an extension to his family including data-related criteria for TIOSTS models. We refine the relation between All-Clock-Resets and the class of Transition-Based Coverage criteria. In his family, All-Paths $\Rightarrow$ All-Clock-Resets, but we prove that the narrow relation All-Transitions $\Rightarrow$ All-Clock-Resets is true too. We introduce the relation All-States $\Rightarrow$ All-Locations that was missing. En-Nouaary's family has neither the All-One-Loop-Paths criterion nor the class

of Data-Flow-Oriented Coverage criteria. We introduce them below the All-Paths criterion. Conversely, our family does not have the All-Clock-Regions criterion, because TIOSTS uses zones instead of regions. Finally, only we evaluate empirically the failure detection capability of eight criteria.

Zhu, Hall and May [32] surveys the literature for test selection criteria at source code level. They present several criteria applicable to unit testing, compare them using the strict inclusion relation and provide an axiomatic study of the properties of criteria. Our work is close to theirs because we also compare test selection criteria using the strict inclusion relation. But we work at model level instead of source code level, and we also perform an empirical study to compare selected criteria.

6 Concluding Remarks

In this paper we presented test selection criteria that can be applied to symbolic transition models of real-time systems, particularly, the TIOSTS model.

We investigated the literature for test selection criteria applicable to models of real-time systems. Next we selected the ones applicable to TIOSTS and formalized a family of 19 test selection criteria partially ordered by the strict inclusion relation.

We evaluated 8 criteria in an empirical study with six TIOSTS models. Our results showed that, even though there are differences on the criteria related to size and failure detection capability, the differences were not significant, particularly when considering time-related failures and cost-effectiveness measured as the rate of size by the number of failures.

In general, we can observe that current specific available criteria are still imprecise, because a number of failures were missed. General criteria were precise, but test suites were large, with a high percentage of test cases that did not fail. Therefore, we can conclude that more effective criteria for the model-based testing of real-time systems are still needed, particularly for symbolic models such as TIOSTS.

As future works, we plan to extend this study to include more test selection criteria, specially the Control-Flow-Oriented Criteria which exercise data and time guards thoroughly. Based on the analysis of advantages and weakness of the criteria in a new empirical study, we intend to propose more precise and effective criteria for TIOSTS.

Acknowledgements. This work was supported by the National Council for Scientific and Technological Development (CNPq) under grants 475710/2013-4, 484643/2011-8, and 560014/2010-4. This work was partially supported by the National Institute of Science and Technology for Software Engineering[9] of CNPq under grant 573964/2008-4. First author was also supported by CNPq. Finally, we thank the anonymous reviewers for their constructive comments.

[9] www.ines.org.br

References

1. Alagar, V.S., Ormandjieva, O., Zheng, M.: Specification-based testing for real-time reactive systems. In: Proceedings of the 34th International Conference on Technology of Object-Oriented Languages and Systems. pp. 25–36 (2000)
2. Almeida, D.R.: Critérios de Geração de Casos de Teste de Sistemas de Tempo Real. Master's thesis, Federal University of Campina Grande, Campina Grande, PB, Brazil (2012)
3. Alur, R., Dill, D.L.: A theory of timed automata. Theoretical Computer Science 126(2), 183–235 (1994)
4. Andrade, W.L., Machado, P.D.L.: Testing interruptions in reactive systems. Formal Aspects of Computing 24, 331–353 (2012)
5. Andrade, W.L., Machado, P.D.L.: Generating test cases for real-time systems based on symbolic models. IEEE Transactions on Software Engineering 39(9), 1216–1229 (2013)
6. Andrade, W.L., Machado, P.D.L., Jéron, T., Marchand, H.: Abstracting time and data for conformance testing of real-time systems. In: Proceedings of the 8th Workshop on Advances in Model Based Testing. pp. 9–17 (2011)
7. Arcuri, A., Iqbal, M.Z., Briand, L.: Black-box system testing of real-time embedded systems using random and search-based testing. In: Proceedings of the 22nd International Conference on Testing Software and Systems. pp. 95–110 (2010)
8. Bengtsson, J., Griffioen, W.O.D., Kristoffersen, K.J., Larsen, K.G., Larsson, F., Pettersson, P., Yi, W.: Verification of an audio protocol with bus collision using UPPAAL. In: Proceedings of the 8th International Conference on Computer Aided Verification. pp. 244–256 (1996)
9. Clarke, D., Lee, I.: Automatic test generation for the analysis of a real-time system: Case study. In: Proceedings of the 3rd IEEE Real-Time Technology and Applications Symposium. pp. 112–124 (1997)
10. El-Far, I.K., Whittaker, J.A.: Model-based software testing. In: Marciniak, J.J. (ed.) Encyclopedia of Software Engineering, vol. 1, pp. 825–837. John Wiley & Sons, Inc. (2002)
11. En-Nouaary, A., Khendek, F., Dssouli, R.: Fault coverage in testing real-time systems. In: Proceedings of the 6th Real-Time Computing Systems and Applications. pp. 150–157 (1999)
12. En-Nouaary, A.: Test selection criteria for real-time systems modeled as timed input-output automata. International Journal of Web Information Systems 3(4), 279–292 (2007)
13. En-Nouaary, A., Hamou-Lhadj, A.: A boundary checking technique for testing real-time systems modeled as timed input output automata. In: Proceedings of the 8th International Conference on Quality Software. pp. 209–215 (2008)
14. Hessel, A.: Model-Based Test Case Selection and Generation for Real-Time Systems. Ph.D. thesis, Uppsala University, Uppsala, Sweden (2007)
15. Jeannet, B., Jéron, T., Rusu, V., Zinovieva, E.: Symbolic test selection based on approximate analysis. In: Proceedings of the 11th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 349–364 (2005)
16. Krichen, M., Tripakis, S.: Black-box conformance testing for real-time systems. In: Proceedings of the 11th International SPIN Workshop on Model Checking of Software. pp. 109–126 (2004)

17. Krichen, M., Tripakis, S.: Conformance testing for real-time systems. Formal Methods in System Design 34(3), 238–304 (2009)
18. Laplante, P.A.: Real-Time System Design and Analysis. John Wiley & Sons (2004)
19. Locke, C.D., Vogel, D.R., Lucas, L., Goodenough, J.B.: Generic avionics software specification. Tech. rep., Software Engineering Institute, Carnegie Mellon University (1990)
20. Nielsen, B., Skou, A.: Test generation for time critical systems: Tool and case study. In: Proceedings of the 13th Euromicro Conference on Real-Time Systems. pp. 155–162 (2001)
21. Peleska, J.: Industrial-strength model-based testing - state of the art and current challenges. In: Proceedings of the 8th Workshop on Model-Based Testing. pp. 3–28 (2013)
22. Pretschner, A., Slotosch, O., Aiglstorfer, E., Kriebel, S.: Model-based testing for real. International Journal on Software Tools for Technology Transfer 5(2), 140–157 (2004)
23. Rapps, S., Weyuker, E.J.: Selecting software test data using data flow information. IEEE Transactions on Software Engineering 11(4), 367–375 (1985)
24. Rusu, V., du Bousquet, L., Jéron, T.: An approach to symbolic test generation. In: Proceedings of the 2nd International Conference on Integrated Formal Methods. pp. 338–357 (2000)
25. Sommerville, I.: Software Engineering. International Computer Science Series, Addison-Wesley, Boston, MA, USA, 9 edn. (2010)
26. Trab, M.S.A., Alrouh, B., Counsell, S., Hierons, R.M., Ghinea, G.: A multi-criteria decision making framework for real time model-based testing. In: Proceedings of the 5th International Academic and Industrial Conference on Testing - Practice and Research Techniques. pp. 194–197 (2010)
27. Tretmans, J.: Model-based testing and some steps towards test-based modelling. In: Proceedings of 11th International School on Formal Methods for the Design of Computer, Communication and Software Systems. pp. 297–326 (2011)
28. Tretmans, J.: Testing concurrent systems: A formal approach. In: Proceedings of the the 10th International Conference on Concurrency Theory. pp. 46–65 (1999)
29. Utting, M., Legeard, B.: Practical Model Based Testing: A Tools Approach. Elsevier, San Francisco, CA, USA (2007)
30. Wohlin, C., Runeson, P., Höst, M., Ohlsson, M.C., Regnell, B., Wesslén, A.: Experimentation in Software Engineering. Springer, New York, NY, USA (2012)
31. Zheng, M., Alagar, V., Ormandjieva, O.: Automated generation of test suites from formal specifications of real-time reactive systems. Journal of Systems and Software 81(2), 286–304 (2008)
32. Zhu, H., Hall, P.A.V., May, J.H.R.: Software unit test coverage and adequacy. ACM Computing Surveys 29(4), 366–427 (1997)

Use Case Analysis based on Formal Methods: An Empirical Study

Marcos Oliveira Junior, Leila Ribeiro, Érika Cota,
Lucio Mauro Duarte, Ingrid Nunes, and Filipe Reis *

PPGC - Institute of Informatics – Federal University of Rio Grande do Sul (UFRGS)
PO Box 15.064 – 91.501-970 – Porto Alegre – RS – Brazil
{marcos.oliveira,leila,erika,lmduarte,ingridnunes,freis}@inf.ufrgs.br

Abstract. *Use Cases (UC)* are a popular way of describing system behavior and UC quality impacts the overall system quality. However, they are presented in natural language, which is usually the cause of issues related to imprecision, ambiguity, and incompleteness. We present the results of an empirical study on the formalization of UCs as Graph Transformation models (GTs) with the goal of running tool-supported analyses on them and revealing possible errors. To evaluate our approach, we apply it to a set of real UC descriptions obtained from a software developer company and measured the results through metrics. The final results demonstrate that this approach can reveal real problems that could otherwise go undetected and, thus, help improve the quality of the UCs.
Keywords. Use Cases, Graph Transformation, Model Analysis.

1 Introduction

Use Cases (UC) [2] are a popular model for documenting software expected behaviour. In current practice, UC descriptions are typically informally documented using, in most cases, natural language in a predefined structure. Being informal descriptions, UCs might be ambiguous and imprecise. Thus, the verification of UCs normally corresponds to manual inspections and walkthroughs [4], and detecting problems is not a trivial task. Since software quality is highly dependent on the quality of the specification, cost-effective strategies to decrease the number of errors in UCs are crucial. Strategies for the formalization of UCs have already been proposed, however, many of them assume a particular syntax for UC description tailored for their particular formalisms. This limits the expression of requirements and, in some cases, also restrains the semantics of the UC. Our aim is to keep the expressiveness of a description in natural language and use a formalism for modeling/analysing UCs that is flexible enough to maintain the semantics defined by stakeholders.

In this paper, we investigate the suitability of Graph Transformation (GT) as a formal model to describe and analyze UCs. Some reasons for choosing GT are: the elements of a UC can be naturally represented as graphs; it is a visual

* This work is partially supported by the VeriTeS project (FAPERGS and CNPq).

language; the semantics is very simple yet expressive; GT is data-driven; there are various static and dynamic analysis techniques available for GT, as well as tools to support them (e.g., [10]). We work towards an approach that integrates UC formalization and tool-supported analysis, with the objective of improving the quality of UCs. We applied our approach on a set of real UC descriptions obtained from a software development company and measured the results.

This paper is organized as follows: Section 2 presents the necessary background information and a overview of the translation of UCs for GTs.; Section 3 presents the settings of the conducted empirical study; Section 4 presents an analysis and discussion of results; Section 5 presents an analysis of threats to the study; Section 6 presents a comparative analysis of our technique with some related work; and Section 7 concludes the paper and discusses future work.

2 Modeling UCs using GTs

2.1 Background

Use Cases. According to Cockburn (2000) [2], a *Use Case (UC)* defines a contract between stakeholders of a system, describing part of the system behavior. The main purpose of a UC description is the documentation of the expected system behavior so as to ease the communication between stakeholders, often including non-technical people, about required system functionalities. For this reason, UC descriptions are usually described in a textual form.
Graph Transformations. The formalism of *Graph Transformations (GT)* [7] is based on defining states of a system as graphs and state changes as rules that transform these graphs. Our analysis of GTs is based on concurrent rules and critical pairs, two methods of analysis independent from the initial state of the system and, thus, they are complementary to any other verification strategy based on initial states (such as testing).

2.2 Proposed Formalization and Verification Approach

The proposed approach, detailed in [6], takes as input a textual UC description, from which the entities and actions that will be part of the formal model are identified. Then, basic verifications can be performed regarding the consistency of the extracted information. If inconsistencies are detected, the UC must be rewritten to eliminate them or the analyst can annotate the problem to be resolved later on. When no basic inconsistencies are found, the GT can then be generated. In this process, conditions and effects of actions are modeled as states and a type graph is built. After that, each UC step is modeled as a transition rule from one state (graph) to another. Having the GT, a series of automatic verifications can be performed to detect possible problems.

We use the AGG tool [10] to perform the automatic analyses on the GT model. All detected issues are annotated as *open issues (OIs)* along with the solutions (when applicable). Open issues are classified according to their severity

level: Yellow (for warnings), Orange (for relevant issues), or Red (for critical issues). The actions to be taken regarding found OIs depend on the analysts, who can determine whether an OI is in fact a real problem.

3 Empirical Study Settings

In order to adequately evaluate our approach, we followed the principles of Experimental Software Engineering [11] and the GQM template [1]. Our main study goal was to demonstrate the usefulness of GTs to improve the quality of UCs by the identification of OIs, from a perspective of the researcher, in the context of a single real software development project. From this, we derived two research questions, which we aimed to answer with our study.

RQ-1 Are system analysts able to detect problems in their own UC descriptions without additional support?

RQ-2 How effective is our GT-based approach in identifying problems in UCs?

The UC descriptions we used in our study are part of the analysis documentation of an industrial software project. This project involves the development of a typical system to manage products from a warehouse, with functional requirements such as adding new products, creating sale orders, and releasing products in stock. We do not provide any further details about our target system and its UCs due to a confidentiality agreement.

3.1 Procedure

The procedure of the study consists of the following steps:

(1) Analysis by System Analyst. We requested a system analyst responsible for the UC descriptions to carefully revise them, and point out problems, such as ambiguity, imprecision, omission, incompleteness, and inconsistency.

(2) UC Formalization. Given a set of UCs, we performed the steps detailed in [6] to formalize them using GTs and used the AGG tool to analyze them, detecting some OIs.

(3) Evaluation of Open Issues. After identifying OIs, we had evaluated whether detected OIs were real problems in the analyzed UCs.

(4) Data Analysis. Our aim is that our approach detects all and only real problems. This can be seen as a *classification problem*, and thus the effectiveness of our approach can be measured using the metrics, widely used in the context of information retrieval, of *precision* and *relative recall* [5], whose formulas are shown below, where *true positives* are OIs that correspond to real problems; $false\ positives$ are OIs that are not real problems; and $false\ negatives$ are real problems not identified as OIs.

$$Precision = \frac{true\ positives}{true\ positives + false\ positives} \tag{1}$$

$$RelativeRecall = \frac{true\ positives}{true\ positives + false\ negatives} \tag{2}$$

Table 1: Study results

OI Type	UC 1		UC 2		UC 3		UC 4		UC 5		Total	
	#OI	#P	#OI	#P	#OI	#P	#OI	#P	#OI	#P	#OI	#P
⚠	3	2	4	2	2	1	4	2	1	0	14	7
?	1	1	1	1	1	1	0	0	2	2	5	5
✋	3	3	1	1	3	2	3	3	3	3	13	12
Total	**7**	**6**	**6**	**4**	**6**	**4**	**7**	**5**	**6**	**5**	**32**	**24**
Legend: UC - Use Case; OI - Open Issue; P - Problem.												

4 Results and Discussion

After revising the original UCs, the system analyst found no problems. However, after applying our approach to these UCs, we identified 32 OIs across the 5 UCs, which gives an average of 6.4 OIs per UC. This is an expressive number, given that the system analyst stated that the UCs had been correctly specified. In order to verify whether the identified OIs were false alarms (false positives), the system analyst was asked to check each one of them. Of the 32 OIs, 24 were pointed out as real problems and only 8 as false positives.

Table 1 presents our results in detail. It shows the number of OIs found in each UC (columns labeled with OI) and how many of these OIs were confirmed as real problems (columns labeled with P). The rows show the number of detected OIs with respect to their level of severity. The table also presents the total number of detected OIs and the total number of real problems considering all the 5 UCs. The symbols ⚠, ?, and ✋ represent warnings (severity Yellow), relevant issues (severity Orange), and critical issues (severity Red), respectively.

We then analyzed these results according to the selected metrics. Because the system analyst was unable to identify any problem without support, the number of problems not identified by our approach was 0, leading to *relative recall* = 1.0. As for the Precision, we obtained 0.75 (24 true positives and 8 false positives) — that is, 75% of the OIs identified by our GT-based approach were real problems. Not only most of the identified issues were actual problems, but also most of the false alarms (7 of 8) were related to low severity OIs.

By analyzing OIs not identified as problems, we observed that 6 of them were not necessarily classified as a false positive by the system analyst. They preferred to leave such issues as they were and postpone changes to future design decisions, considering that they alone could not decide what was the best approach to tackle those issues. The other 2 OIs found, confirmed as false positives, were related to incompleteness or ambiguities due to the lack of knowledge of the modeler about the problem domain and the internal processes of the company.

Note that OIs were identified without the intervention of any stakeholder. The only provided input was the software documentation in the form of UC descriptions and the output was a checklist with OIs to be revised. More importantly, had these problems been detected before the design and implementation, when they should have, development costs could have been potentially reduced.

5 Threats to Validity

During our study we carefully considered validity concerns. This section discusses the main threats we identified to validate this study and how we mitigated them.

Internal Validity. The main threat to internal validity of this study was the selection of modeler of UCs in the formalism of graphs. However, we want to show that, correctly following the steps of our strategy, the modeler does not need a deep understanding of the formalism. Moreover, we used the AGG tool to automate the analyses and provide a graphical interface.

Construct Validity. There are different ways of modeling a system through the formalism of graphs that can produce some threats to construct validity. The modeler may not follow correctly the modeling steps, being influenced by their prior knowledge about the formalism. Therefore, we proposed a roadmap, step by step, on how to model UCs as GTs, for both beginners and experts users.

Conclusion Validity. As the main threat to validity of the conclusion we highlight potential problems in the generation of the model in the formalism of graphs. Once again, our step-by-step modeling process should be followed to prevent the modeler from creating a model that is not consistent with the textual description. Moreover, the tool-supported verifications can also detect such modeling errors, thus reducing the risk of this threat.

External Validity. The main threat to the external validity was the selection of artifacts on which we based our study. We did not use any criteria to select either the project or the system analyst who participated of our study. We were aware of this threat during the study. However, we opted for randomly choosing artifacts to support the applicability of our strategy in different scenarios.

6 Related Work

Some authors have developed approaches for translating UCs to well-known formalisms, such as LTS [8], Petri Nets [12], and FSM [9]. Unlike these formalisms, a GT model is data-driven and we do not need to explicitly determine the control flow unless it is necessary to guarantee data consistency. The approach presented in [13] allows the simulation of the execution of the system but do not report the use of any type of analysis, which, in our opinion, reduces the advantage of having a formal model. The work described in [3] considers analyses such as critical pairs and dependencies involving multiple UCs and provides some ideas on the interpretation of the results. However, we propose a more structured way of providing diagnostic feedback about single UCs, which serves as a guide to point out the possible errors as well as their severity level.

7 Conclusions and Future Work

We investigated the suitability of GT as a formal basis for UC description and improvement. We evaluated our approach through an experiment with real software artifacts, where we could detect existing errors, which helped improve the

original UCs. Making a general analysis of the experiment, we consider the results promising, since it was possible to identify a large number of real problems based on a documentation that was produced at an early stage of software development. We observed the need for further automating the process, if not all, at least some steps, which is one of the most immediate planned future work.

A inter-UC analysis is currently being implemented as well as a more detailed diagnostic feedback. Within the same model frame, other types of validation and verification techniques on GT models, such as test case generation, model checking, and theorem proving, are also subject of current work. We plan to investigate whether we could reduce the impact and cost of changes by identifying which parts of the description are affected. Finally, note that, although we did not present any new formal method or verification technique here, a considerable amount of expertise in formal methods was required to define the OIs: they are meant to bridge the gap between the informal and formal worlds. We believe that this type of work is crucial towards the industrial adoption of formal methods.

References

1. Basili, V., Caldiera, C., Rombach, H.: Goal Question Metric Paradigm, Encyclopedia of Software Engineering, vol. 1. John Wiley & Sons (1994)
2. Cockburn, A.: Writing Effective Use Cases. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1st edn. (2000)
3. Hausmann, J.H., Heckel, R., Taentzer, G.: Detection of conflicting functional requirements in a use case-driven approach: A static analysis technique based on graph transformation. In: Proc. of the 24th ICSE. pp. 105–115 (2002)
4. Myers, G., Sandler, C., Badgett, T.: The Art of Software Testing. ITPro Collection, Wiley (2011)
5. Powers, D.M.: Evaluation: From Precision, Recall and F-Factor to ROC, Informedness, Markedness & Correlation. Tech. Rep. SIE-07-001, Flinders University of South Australia (2007)
6. Ribeiro, L., Cota, E., Duarte, L.M., Oliveira Jr., M.A.d.: Improving the quality of use cases via model construction and analysis. In: Proc. of the 22nd WADT (2014)
7. Rozenberg, G. (ed.): Handbook of graph grammars and computing by graph transformation: volume I: Foundations. World Scientific, River Edge, USA (1997)
8. Sinnig, D., Chalin, P., Khendek, F.: LTS semantics for use case models. In: Proc. of the ACM SAC. pp. 365–370. ACM (2009)
9. Sinnig, D., Chalin, P., Khendek, F.: Use case and task models: An integrated development methodology and its formal foundation. ACM ToSEM 22(3), 27:1–27:31 (Jul 2013)
10. Taentzer, G.: AGG: A tool environment for algebraic graph transformation. In: Applications of Graph Transformations with Industrial Relevance, LNCS, vol. 1779, pp. 481–488. Springer Berlin Heidelberg (2000)
11. Wohlin, C., Runeson, P., Höst, M., Ohlsson, M.C., Regnell, B., Wesslén, A.: Experimentation in Software Engineering. Springer Berlin Heidelberg (2012)
12. Zhao, J., Duan, Z.: Verification of use case with petri nets in requirement analysis. In: Proc. of the ICCSA: Part II. pp. 29–42. Springer-Verlag (2009)
13. Ziemann, P., Hǎvlscher, K., Gogolla, M.: From UML models to graph transformation systems. ENTCS 127(4), 17 – 33 (2005)

A dynamic logic for every season

Alexandre Madeira[1], Renato Neves[1], Manuel A. Martins[2] and Luís S. Barbosa[1]

HASLab INESC TEC & Univ. Minho
{madeira@di.uminho.pt, nevrenato@di.uminho.pt, lsb@di.uminho.pt}
CIDMA - Dep. Mathematics, Univ. Aveiro
martins@ua.pt

Abstract. This paper introduces a method to build dynamic logics with a graded semantics. The construction is parametrized by a structure to support both the spaces of truth and of the domain of computations. Possible instantiations of the method range from classical (assertional) dynamic logic to less common graded logics suitable to deal with programs whose transitional semantics exhibits fuzzy or weighted behaviour. This leads to the systematic derivation of program logics tailored to specific program classes

1 Introduction

Propositions, capturing static properties of program states, and events, or actions, which are responsible for transitions from a state to another, are the key ingredients in modelling and reasoning about state-based software systems. The latter are typically combined through a Kleene algebra to express sequential, non deterministic, iterative behaviour of systems, while the former brings to the scene a logical structure.

Dynamic logic [6], a generalisation of the logic of Floyd-Hoare, is a well known and particularly powerful way of combining these two dimensions into a formal framework to reason about computational systems. Its potential stems from blending together classical logic, enriched with a modal dimension to express system's dynamics, and a (Kleene) algebra of actions to structure programs.

Over time dynamic logic grew to an entire family of logics increasingly popular in the verification of computational systems, and able to evolve and adapt to new, and complex validation challenges. One could mention its role in model validation (as in *e.g.* [10]), or the whole family of variants tailored to specific programming languages (as in *e.g.* [11,1]), or its important extensions to new computing domains, namely probabilistic [8] or continuous [13,14].

The latter is particularly relevant from an Engineering point of view: Actually, Platzer's hybrid dynamic logic, and its associated tool, KeYmaera, combining an algebra of actions based on real numbers assignments, the standard Kleene operators and differential equations to specify continuous transitions from the "real" (physical) world, provides a powerful framework for the design and validation of cyber-physical systems with increased industrial relevance [16].

If cyber-physical systems gives rise to the need for ways of dealing with continuous state spaces, in a number of other cases dealing with some form of "quantitative" transitions (weighted, probabilistic, etc) is also a must. Hence the quest for dynamic logics able to capture smoothly these kind of phenomena is becoming more and more important.

This paper intends to contribute in this path. In particular, our attention is focussed on graded logics [4,17], in the broad sense of attaching partially ordered grades to logical formulas to express, in one way or another, uncertain information. In this broad sense, fuzzy [5], probabilistic [12] or weighted logics [2] may be brought into the picture.

In this context, the purpose of this work is the development of a generic method to construct graded dynamic logics. Technically, the definition of these logics is parametrized by a (specific kind of) an action lattice [7] which combines a (slight generalisation of a) Kleene algebra with a residuated lattice structure. The latter captures the graded logic dimension and fits nicely with our objectives. Moreover, the extension of Kleene algebras with residuation operators, providing weak right and left inverses to sequential composition as in [15], as well as with a lattice structure leads to a finitely-based equational variety which, as plain Kleene algebras, is closed under the formation of square matrices [9].

The relevance of this closure property lies in the fact that several problems modelled as (weighted) transition systems can be formulated as matrices over a Kleene algebra or a related structure. Following such a trend, we represent programs as matrices supporting the information about their effects when executed from each state in the state space. The interested reader is referred to [3] for a detailed discussion on the relationship between Kleene algebras, action algebras and action lattices.

The remaining of this paper is organised as follows. Section 2 recalls from [7] the definition of an action lattice and introduces a method, parametric on such a lattice, to generate graded dynamic logics. The construction put forward is illustrated with several examples. Then, in section 3, it is shown that the resulting logic is a dynamic logic indeed, in the sense that all the rules of propositional dynamic logic restricted to positive-existential formulas still hold. Finally, section 4 concludes and suggests points for future research.

2 The method

This section introduces a generic method to generate *graded dynamic logics* parametric on a complete action lattice which captures both the structure of the computational domain and that of the (logical) truth space.

Let us start by recalling from [7] the following definition:

Definition 1. *An* action lattice *is a tuple*

$$\mathbf{A} = (A, +, ;, 0, 1, *, \leftarrow, \rightarrow, \cdot)$$

where, for A a set, 0 and 1 are constants and $+, ;, *, \leftarrow, \rightarrow$ and $\cdot$ are binary operations in A satisfying the axioms in Figure 1. Relation $\leq$ is the one induced by $+$ as $a \leq b$ iff $a + b = b$.

$$
\begin{aligned}
a + (b + c) &= (a + b) + c && (1)\\
a + b &= b + a && (2)\\
a + a &= a && (3)\\
a + 0 &= 0 + a = a && (4)\\
a; (b; c) &= (a; b); c && (5)\\
a; 1 &= 1; a = a && (6)\\
a; (b + c) &= a; b + a; c && (7)\\
(a + b); c &= a; c + b; c && (8)\\
a; 0 &= 0; a = 0 && (9)\\
1 + a + a^*; a^* &\leq a^* && (10)\\
a; x \leq x &\Rightarrow a^*; x \leq x && (11)\\
x; a \leq x &\Rightarrow x; a^* \leq x && (12)\\
a; x \leq b &\Leftrightarrow x \leq a \rightarrow b && (13)\\
x; a \leq b &\Leftrightarrow x \leq a \leftarrow b && (14)\\
(x \rightarrow x)^* &= x \rightarrow x && (15)\\
(x \leftarrow x)^* &= x \leftarrow x && (16)\\
a \cdot (b \cdot c) &= (a \cdot b) \cdot c && (17)\\
a \cdot b &= b \cdot a && (18)\\
a \cdot a &= a && (19)\\
a + (a \cdot b) &= a && (20)\\
a \cdot (a + b) &= a && (21)
\end{aligned}
$$

Fig. 1. Axiomatisation of action lattices (from [7])

An action lattice is said to be complete when equipped with both a supremum and an infimum of all subsets of A. Therefore, complete action lattices have biggest and smallest elements denoted in the sequel by $\top$ and $\bot$, respectively. Note that in any action lattice, $\bot = 0$, since for any $a \in A$, $a + 0 = a$, i.e., $0 \leq a$. In this paper we resort to notation $\sum$ for the iterated version of the (join) operator $+$, and to notation $\prod$ for the iterated version of the (meet) operator $\cdot$.

The starting point for the method proposed here is thus the choice of an appropriate action lattice

$$\mathbf{A} = (A, +, ;, 0, 1, *, \leftarrow, \rightarrow, \cdot)$$

Additionally, we require $\mathbf{A}$ to satisfy the following equality:

$$a;(b \cdot c) = a;b \cdot a;c \tag{22}$$

As mentioned above, this structure supports both the computational paradigm (to distinguish between *e.g.* imperative, deterministic or non deterministic computations, or between plain or weighted transitions) and the truth space (to capture *e.g.* the standard Boolean reasoning or more complex truth spaces). Before proceeding let us exemplify this structure with a couple of action lattices typically found in Computer Science applications. In the examples, the logic generated by an action lattice $\mathbf{A}$ will be denoted by $\mathcal{GDL}(\mathbf{A})$.

Example 1 ($\mathcal{GDL}(\mathbf{2})$ — the standard propositional dynamic logic.). Standard propositional dynamic logic is generated from the following structure

$$\mathbf{2} = (\{\top, \bot\}, \vee, \wedge, \bot, \top, *, \leftarrow, \rightarrow, \wedge)$$

with the standard boolean connectives:

$\vee$	$\bot$	$\top$
$\bot$	$\bot$	$\top$
$\top$	$\top$	$\top$

$\wedge$	$\bot$	$\top$
$\bot$	$\bot$	$\bot$
$\top$	$\bot$	$\top$

$\rightarrow$	$\bot$	$\top$
$\bot$	$\top$	$\top$
$\top$	$\bot$	$\top$

$*$	
$\bot$	$\top$
$\top$	$\top$

and taking $a \leftarrow b$ whenever $b \rightarrow a$. It is not difficult to see that $\mathbf{2}$ is an action algebra. Moreover, the lattice is obviously complete and it satisfy the condition (22) (note that both composition and the meet operator are realized by $\wedge$).

Example 2 ($\mathcal{GDL}(\mathbf{3})$ — a dynamic logic to deal with unknown data.). This is a three-valued logic, with an explicit representative for *unknown*, or *uncertain* information. Note that the three elements linear lattice induces an action lattice

$$\mathbf{3} = (\{\top, u, \bot\}, \vee, \wedge, \bot, \top, *, \leftarrow, , \rightarrow, \wedge)$$

where

$\vee$	$\bot$	u	$\top$
$\bot$	$\bot$	u	$\top$
u	u	u	$\top$
$\top$	$\top$	$\top$	$\top$

$\wedge$	$\bot$	u	$\top$
$\bot$	$\bot$	$\bot$	$\bot$
u	$\bot$	u	u
$\top$	$\bot$	u	$\top$

$\rightarrow$	$\bot$	u	$\top$
$\bot$	$\top$	$\top$	$\top$
u	u	$\top$	$\top$
$\top$	$\bot$	u	$\top$

$*$	
$\bot$	$\top$
u	$\top$
$\top$	$\top$

and taking $a \leftarrow b$ whenever $b \rightarrow a$. It is easy to see all the conditions in Definition 1 hold. Moreover, the lattice is obviously complete and satisfies condition (22). The refer should note that both composition and meet are realized by $\wedge$).

Example 3 ($\mathcal{GDL}(\text{Ł})$ — a dynamic logic to deal with continuous levels of fuzziness.).

This is based on the well-known Łukasiewicz arithmetic lattice

$$\text{Ł} = ([0,1], max, min, 0, 1, *, \rightarrow, \leftarrow, min)$$

where

– $x \to y \;=\; min\{1, 1 - x + y\}$,
– $x \leftarrow y \;=\; 1 - max\{0, x + y - 1)\}$ and
– $*$ maps each point of $[0, 1]$ to 1.

Again, this defines a complete action lattice which additionally satisfies condition (22). Note that both composition and the meet operator are now represented by function min.

Example 4 ($\mathcal{GDL}$(FW) – a dynamic logic to deal with resource consuming systems.). This example explores the so called Floyd-Warshall algebra which consists of a tuple

$$\mathbb{N}^{+}_{\bot\top} = (\{\bot, 0, 1, \ldots, \top\}, max, +, \bot, 0, *, \smile, \smile, min)$$

where $\smile$ is the truncated subtraction

$$a \smile b = \begin{cases} a - b, & \text{if } a \geq b \\ \bot & \text{otherwise} \end{cases}$$

and, for any $i > 0$,

$*$	
$\bot$	0
0	0
i	$\top$

The induced order $\leq$ corresponds to $\leq$ in $\mathbb{N}$ ($a \leq b$ iff $max\{a, b\} = b$). The lattice is also complete and it satisfies condition (22) because $a + min\{b, c\} = min\{a + b, a + c\}$.

Illustrated the notion of an action lattice, we are now prepared to introduce the general construction of graded dynamic logics. We consider now its signatures, formulæ, semantics and satisfaction. Thus,

Signatures. Signatures of $\mathcal{GDL}(\mathbf{A})$ are pairs (Π, Prop) corresponding to the denotations of atomic computations and of propositions, respectively.

Formulæ. A core ingredient of any dynamic logic is its set of programs. Therefore, let us denote the set of atomic programs by Π. The *set of Π-programs*, denoted by $Prog(\Pi)$, consists of all the expressions generated by

$$\pi \ni \pi_0 \,|\, \pi; \pi \,|\, \pi + \pi \,|\, \pi^*$$

for $\pi_0 \in \Pi$. Given a signature (Π, Prop), we define the $\mathcal{GDL}(\mathbf{A})$-formulas for (Π, Prop), denoted by $\mathrm{Fm}^{\mathcal{GDL}(\mathbf{A})}(\Pi, \mathrm{Prop})$, by the grammar

$$\rho \ni \top \,|\, \bot \,|\, p \,|\, \rho \vee \rho \,|\, \rho \wedge \rho \,|\, \rho \to \rho \,|\, \langle \pi \rangle \rho$$

for $p \in \mathrm{Prop}$ and $\pi \in Prog(\Pi)$. Note that this corresponds to the *positive existential* fragment of the propositional dynamic logic.

Semantics. The first step is to introduce the space where the computations of $\mathcal{GDL}(\mathbf{A})$ are to be interpreted. As usual, this corresponds to a Kleene algebra. Therefore, we consider the structure

$$\mathbb{M}_n(\mathbf{A}) = (M_n(\mathbf{A}), +, ;, \mathbf{0}, \mathbf{1}, ^*)$$

defined as follows:

1. $M_n(\mathbf{A})$ is the space of the $(n \times n)$-action lattices over $\mathbf{A}$
2. for any $A, B \in M_n(Q)$, define $M = A+B$ by $M_{i,j} = A_{i,j} + B_{i,j}$, $i, j \leq n$.
3. for any $A, B \in M_n(Q)$, define $M = A \,;\, B$ by taking $M_{i,j} = \sum_{k=1}^{n}(A_{i,k}; B_{k,j})$ for any $i, j \leq n$.
4. $\mathbf{1}$ and $\mathbf{0}$ are the $(n \times n)$-matrices defined by $\mathbf{1}_{i,j} = 1$ and $\mathbf{0}_{i,j} = 0$, for any $i, j \leq n$.
5. for any $M = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right] \in M_n(Q)$, where A and D are square matrices, we take

$$M^* = \left[\begin{array}{c|c} F^* & F^* \,;\, B \,;\, D^* \\ \hline C \,;\, F^* & D^* + D^* \,;\, C \,;\, F^* \,;\, B \,;\, D^* \end{array}\right]$$

   where $F = A + B \,;\, D^* \,;\, C$. Note that this construction is recursively defined from the base case (where $n = 2$) where the operations of the base action lattice $\mathbf{A}$ are used.

Finally, we have to show that,

Theorem 1. *The structure* $\mathbb{M}_n(\mathbf{A}) = (M_n(\mathbf{A}), +, ;, \mathbf{0}, \mathbf{1}, ^*)$ *defined as above is a Kleene algebra.*

Proof. The structure, and the respective operations, corresponds to the algebra of matrices over $(A, +, ;, 0, 1, *)$, *i.e.*, the Kleene algebra underlying action lattice $\mathcal{A}$. A canonical result establishes that Kleene algebras are closed under formation of matrices (e.g. [9]). Therefore, $\mathbb{M}_n(\mathbf{A})$ constitutes a Kleene algebra. □

$\mathcal{GDL}(\mathbf{A})$-models for a set of propositions Prop, denoted by $\mathrm{Mod}^{\mathcal{GDL}(\mathbf{A})}(\mathrm{Prop})$, consists of tuples

$$\mathcal{A} = (W, V, (\mathcal{A}_\pi)_{\pi \in \Pi})$$

where

- W is a finite set (of states),
- $V : \mathrm{Prop} \times W \to A$ is a function,
- and $\mathcal{A}_\pi \in \mathbb{M}_n(\mathbf{A})$, with n standing for the cardinality of W.

The interpretation of programs in these models is made by matrices over the Kleene algebra of $\mathbf{A}$. Each matrix represents the effect of a program executing from any point of the model. Formally, the interpretation of a program $\pi \in Prog(\Pi)$ in a model $\mathcal{A} \in \mathrm{Mod}^{\mathcal{GDL}(\mathbf{A})}(\Pi, \mathrm{Prop})$ is recursively defined, from the atomic programs $(\mathcal{A}_\pi)_{\pi \in \Pi}$, as follows:

– $\mathcal{A}_{\pi;\pi'} = \mathcal{A}_\pi \,;\mathcal{A}_{\pi'}$
– $\mathcal{A}_{\pi+\pi'} = \mathcal{A}_\pi + \mathcal{A}_{\pi'}$
– $\mathcal{A}_{\pi^*} = \mathcal{A}_\pi^*$

Observe that the set of states W supports the index system of the programs (adjacency) matrices. In this context, it is important to note, that, for example,

$$((M_{\pi;\pi'})_{ij} = \sum_{k=1}^{n} \{(M_\pi)_{ik}; (M_{\pi'})_{kj}\}$$

corresponds to

$$((M_{\pi;\pi'}(w, w')) = \sum_{w'' \in W} \{(M_\pi)(w, w''); (M_{\pi'}(w'', w')\}$$

where i and j stands for the adjacency index of w and w', respectively. Actually, the latter characterisation is often used in the sequel.

Example 5 (Computations spaces).

Let us fix an action lattice $\mathbf{A} = (A, +, ;, 0, 1, *, \leftarrow, \rightarrow, \cdot)$ and a signature $(\{\pi, \pi'\}, \{p\})$. Then, consider a model $\mathcal{A} = (W, V, (\mathcal{A}_\pi)_{\pi \in \Pi})$, with $W = \{s_1, s_2\}$ and the following atomic programs

$$\mathcal{A}_\pi = \begin{bmatrix} \bot & q_{12} \\ \bot & q_{22} \end{bmatrix} \qquad \mathcal{A}_{\pi'} = \begin{bmatrix} \bot & q'_{12} \\ \bot & \bot \end{bmatrix}$$

which can be represented by the following labelled transition system:

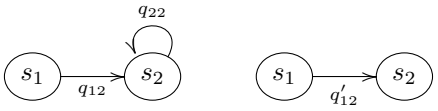

Let us suppose that $\mathcal{A}$ is realized by

$$\mathbf{2} = (\{\top, \bot\}, \vee, \wedge, \bot, \top, *, \leftarrow, \rightarrow, \wedge)$$

Making $q_{12} = q_{22} = q'_{1,2} = \top$ we get the standard adjacency matrices of the graph underlying the transition system. In this case, we interpret choice $\pi + \pi'$ and composition $\pi; \pi'$ by

$$\mathcal{A}_{\pi+\pi'} = \mathcal{A}_\pi + \mathcal{A}_{\pi'} = \begin{bmatrix} \bot & \top \\ \bot & \top \end{bmatrix} + \begin{bmatrix} \bot & \top \\ \bot & \bot \end{bmatrix} = \begin{bmatrix} \bot \vee \bot & \top \vee \top \\ \bot \vee \bot & \top \vee \bot \end{bmatrix} = \begin{bmatrix} \bot & \top \\ \bot & \top \end{bmatrix}$$

The interpretation of the composition $\pi; \pi'$ is computed as follows,

$$\mathcal{A}_{\pi;\pi'} = \begin{bmatrix} \bot & \top \\ \bot & \top \end{bmatrix} ; \begin{bmatrix} \bot & \top \\ \bot & \bot \end{bmatrix} = \begin{bmatrix} (\bot \wedge \bot) \vee (\top \wedge \bot) & (\bot \wedge \top) \vee (\top \wedge \bot) \\ (\bot \wedge \bot) \vee (\top \wedge \bot) & (\bot \wedge \top) \vee (\top \wedge \bot) \end{bmatrix} = \begin{bmatrix} \bot & \bot \\ \bot & \bot \end{bmatrix}$$

The interested reader can easily verify that, as expected,

$$\mathcal{A}_{\pi';\pi} = \begin{bmatrix} \bot & \top \\ \bot & \bot \end{bmatrix}$$

For the interpretation of the π closure, we have

$$\mathcal{A}_{\pi^*} = (\mathcal{A}_\pi)^* \begin{bmatrix} \bot & \top \\ \bot & \top \end{bmatrix}^* = \begin{bmatrix} f^* & f^* \wedge \top \wedge \top^* \\ \bot & \top^* \vee (\top^* \wedge \bot \wedge \top \wedge \top) \end{bmatrix}$$

where $f = \bot \vee (\top \wedge \top^* \wedge \bot) = \bot$, hence

$$\mathcal{A}_{\pi *} = \begin{bmatrix} \top & \top \\ \bot & \top \end{bmatrix}$$

as expected.

The reader is now invited to a small exercise. Taking the same matrix in the context of

$$\mathbf{3} = (\{\top, u, \bot\}, \vee, \wedge, \bot, \top, *, \leftarrow, \rightarrow, \wedge)$$

and considering $q_{12} = q_{22} = \top$ and $q'_{12} = u$, let us compute composition

$$\mathcal{A}_{\pi';\pi} = \begin{bmatrix} \bot & u \\ \bot & \bot \end{bmatrix} ; \begin{bmatrix} \bot & \top \\ \bot & \top \end{bmatrix} = \begin{bmatrix} (\bot \wedge \bot) \vee (u \wedge \bot) & (\bot \wedge \top) \vee (u \wedge \top) \\ (\bot \wedge \bot) \vee (\bot \wedge \top) & (\bot \wedge \top) \vee (\bot \wedge \top) \end{bmatrix} = \begin{bmatrix} \bot & u \\ \bot & \bot \end{bmatrix}$$

As expected, the unknown factor affecting transition $s_1 \rightarrow s_2$ in $\mathcal{A}'_\pi$ is propagated to transition $s_2 \rightarrow s_2$ in $\mathcal{A}_{\pi';\pi}$.

If a continuous space is required to define the "unknown metric", one may resort to the Łukasiewicz arithmetic lattice

$$Ł = ([0,1], max, min, 0, 1, *, \rightarrow, \leftarrow, min)$$

Consider, for instance, $q_{12} = a$, $q_{22} = b$ and $q'_{12} = c$ for some $a, b, c \in [0,1]$. In this case we may, for example, compute choice $\pi + \pi'$, making

$$\mathcal{A}_{\pi+\pi'} = \begin{bmatrix} 0 & a \\ 0 & b \end{bmatrix} + \begin{bmatrix} 0 & c \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} max\{0,0\} & max\{a,c\} \\ max\{0,0\} & max\{b,0\} \end{bmatrix} = \begin{bmatrix} 0 & max\{a,c\} \\ 0 & b \end{bmatrix}$$

The reader may check that

$$\mathcal{A}_{\pi^*} = \begin{bmatrix} 1 & a \\ 0 & 1 \end{bmatrix}$$

Note that the certainty value 1 in the diagonal of the matrix stands for the reflexive dimension of the reflexive-transitive closure $*$.

Let us now consider the action lattice

$$\mathbb{N}^+_{\bot\top} = (\{\bot, 0, 1, \ldots, \top\}, max, +, \bot, 0, *, \backsim, \backsim, min)$$

As stated above, the structure $\mathbb{N}^+_{\bot\top}$ is suitable to reason about resource consuming systems. The value of a transition is $\top$ when it costs an infinite amount

of resources; it is $\bot$ when undefined. The composition of actions reflects the accumulation of sequential costs. For instance $\mathcal{A}_{\pi';\pi} =$

$$\begin{bmatrix} \bot & c \\ \bot & \bot \end{bmatrix} ; \begin{bmatrix} \bot & a \\ \bot & b \end{bmatrix} = \begin{bmatrix} max\{\bot+\bot, c+\bot\} & max\{\bot+a, c+b\} \\ max\{\bot+\bot, \bot+\bot\} & max\{\bot+a, \bot+b\} \end{bmatrix} = \begin{bmatrix} \bot & c+b \\ \bot & \bot \end{bmatrix}$$

Moreover, the interpretation of a program $\pi + \pi'$ reflects, in each transition the most expensive choice:

$$\mathcal{A}_{\pi'+\pi} = \begin{bmatrix} \bot & c \\ \bot & \bot \end{bmatrix} + \begin{bmatrix} \bot & a \\ \bot & b \end{bmatrix} = \begin{bmatrix} max\{\bot, \bot\} & max\{c, a\} \\ max\{\bot, \bot\} & max\{\bot, b\} \end{bmatrix} = \begin{bmatrix} \bot & max\{c, a\} \\ \bot & b \end{bmatrix}$$

Finally, observe the interpretation of the closure of π

$$\mathcal{A}_{\pi^*} = \begin{bmatrix} \bot & a \\ \bot & b \end{bmatrix} = \begin{bmatrix} f^* & f^* + a + b^* \\ \bot & max\{b^*, b^* + \bot + \bot^* + a + b^*\} \end{bmatrix} = \begin{bmatrix} 0 & a + b^* \\ \bot & b^* \end{bmatrix}$$

where $f = max\{\bot, a + b^* + \bot\}$. Note that for any $b > 0$, the matrix assumes

$$\begin{bmatrix} 0 & \top \\ \bot & \top \end{bmatrix}$$

which reflects the cost of an undetermined repetition of transition $s_2 \to s_2$. Naturally, when the cost of the action is 0, we have

$$\begin{bmatrix} 0 & a \\ \bot & 0 \end{bmatrix}$$

Satisfaction. Finally, let us characterise the (graded) satisfaction relation. As mentioned above, the carrier of $\mathbf{A}$ corresponds to the space of truth degrees for $\mathcal{GDL}(\mathbf{A})$. Hence, the graded satisfaction relation for a model $\mathcal{A} \in \mathrm{Mod}^{\mathcal{GDL}(\mathbf{A})}(\Pi, \mathrm{Prop})$ consists of a function

$$\models : \mathrm{Fm}^{\mathcal{GDL}(\mathbf{A})}(\Pi, \mathrm{Prop}) \times W \to A$$

recursively defined as follows:

- $(w \models \top) = \top$
- $(w \models \bot) = \bot$
- $(w \models p) = V(p, w)$, for any $p \in \mathrm{Prop}$
- $(w \models \rho \wedge \rho') = (w \models \rho) \cdot (w \models \rho')$
- $(w \models \rho \vee \rho') = (w \models \rho) + (w \models \rho')$
- $(w \models \rho \to \rho') = (w \models \rho) \to (w \models \rho')$
- $(w \models \langle \pi \rangle \rho) = \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); (w' \models \rho)\}$

Example 6. In order to make a case for the versatility and generality of this method, let us consider the evaluation of the very simple sentence $\langle \pi^* \rangle p$ in three

of the dynamic logics constructed in the examples above. Concretely, let us evaluate $\langle \pi^* \rangle p$ in state s_1. For this we calculate

$$(s_1 \models \langle \pi^* \rangle p) = \sum_{w' \in W} \{\mathcal{A}_{\pi^*}(s_1, w'); (w' \models p)\}$$

Starting with $\mathcal{GDL}(\mathbf{2})$, let us assume $V(p, s_1) = \bot$ and $V(p, s_2) = \top$. In this case, as expected

$$\begin{aligned}(s_1 \models \langle \pi^* \rangle p) &= \textstyle\sum_{w' \in W} \{\mathcal{A}_{\pi^*}(s_1, w'); (w' \models p)\} \\ &= \big(\mathcal{A}_{\pi^*}(s_1, s_1) \wedge (s_1 \models p)\big) \vee \big(\mathcal{A}_{\pi^*}(s_1, s_2) \wedge (s_2 \models p)\big) \\ &= (\top \wedge V(p, s_1)) \vee (\top \wedge V(p, s_2)) \\ &= (\top \wedge \bot) \vee (\top \wedge \top) \\ &= \top\end{aligned}$$

This means that we can achieve p from s_1 through π^*. Considering the $\mathcal{GDL}(\text{Ł})$ and assuming $V(s_1, p) = 0$ and $V(s_2, p) = 1$, we may calculate

$$\begin{aligned}(s_1 \models \langle \pi^* \rangle p) &= \textstyle\sum_{w' \in W} \big\{\mathcal{A}_{\pi^*}(s_1, w'); (w' \models p)\big\} \\ &= max\big\{min\{\mathcal{A}_{\pi^*}(s_1, s_1), (s_1 \models p)\}, min\{\mathcal{A}_{\pi^*}(s_1, s_2), (s_2 \models p)\}\big\} \\ &= max\big\{min\{1, 0\}, min\{a, 1\}\big\} \\ &= max\{0, a\} \\ &= a\end{aligned}$$

Therefore, we can assure, with a degree of certainty a, that we can achieve p from s_1 through π^*.

Interpreting now the same sentence in logic $\mathcal{GDL}(\mathbb{N}^+_{\bot\top})$, assuming that $V(s_1, p) = \bot$ and $V(s_2, p) = 0$, we get

$$\begin{aligned}(s_1 \models \langle \pi^* \rangle p) &= \textstyle\sum_{w' \in W} \big\{\mathcal{A}_{\pi^*}(s_1, w'); (w' \models p)\big\} \\ &= max\big\{\mathcal{A}_{\pi^*}(s_1, s_1) + (s_1 \models p), \mathcal{A}_{\pi^*}(s_1, s_2) + (s_2 \models p)\big\} \\ &= max\{0 + \bot, a + b^* + 0\} \\ &= a + b^*\end{aligned}$$

Hence, we can say that p can be accessed from s_1 through π^* consuming $a + b^*$ resources unities.

3 "Dynamisations" are dynamic

Having introduced a generic method for generating dynamic logics, this section establishes that the resulting logics behave, in fact, as dynamic logics. In particular, all the axioms of the propositional dynamic logic involving positive-existential formulas (see [6]) remain sound in this generic construction.

In the context of graded satisfaction, the verification that a property ρ is valid corresponds to the verification that, for any state w of any model $\mathcal{A}$, $(w \models \rho) = \top$. Hence, by (13) and (14), we have that asserting $(\rho \leftrightarrow \rho') = \top$ is equivalent to prove that, for any $w \in W$, that $(w \models \rho) = (w \models \rho')$; and to proof $(\rho \rightarrow \rho') = \top$ is equivalent to proof that $(w \models \rho) \leq (w \models \rho')$.

Lemma 1. *The following are valid formulas in any* $\mathcal{GDL}(\mathbf{A})$*:*

(1.1) $\langle\pi\rangle(\rho \vee \rho') \leftrightarrow \langle\pi\rangle(\rho) \vee \langle\pi\rangle\rho'$
(1.2) $\langle\pi\rangle(\rho \wedge \rho') \rightarrow \langle\pi\rangle(\rho) \wedge \langle\pi\rangle\rho'$

Proof. **Axiom (1.1)**

$$
\begin{aligned}
& (w \models \langle\pi\rangle(\rho \vee \rho')) \\
= & \quad \{ \text{ defn of } \models \} \\
& \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); (w' \models \rho \vee \rho')\} \\
= & \quad \{ \text{ defn. of } \models \} \\
& \sum_{w' \in W} \{(\mathcal{A}_\pi(w, w'); \big((w' \models \rho) + (w' \models \rho')\big)\} \\
= & \quad \{ (7) \} \\
& \sum_{w' \in W} \{(\mathcal{A}_\pi(w, w'); (w' \models \rho) + (\mathcal{A}_\pi(w, w'); (w' \models \rho'))\} \\
= & \quad \{ \text{ supremum properties} \} \\
& \sum_{w' \in W} \{(\mathcal{A}_\pi(w, w'); (w' \models \rho)\} + \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); (w' \models \rho'))\} \\
= & \quad \{ \text{ defn of } \models \} \\
& (w \models \langle\pi\rangle\rho) + (w \models \langle\pi\rangle\rho) \\
= & \quad \{ \text{ defn of } \models \} \\
& (w \models \langle\pi\rangle\rho \vee \langle\pi\rangle\rho)
\end{aligned}
$$

Therefore $\langle\pi\rangle(\rho \vee \rho') \leftrightarrow \langle\pi\rangle\rho \vee \langle\pi\rangle\rho$.

Axiom (1.2)

$$
\begin{aligned}
& (w \models \langle\pi\rangle(\rho \wedge \rho')) \\
= & \quad \{ \text{ defn of } \models \} \\
& \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); (w' \models \rho \wedge \rho')\} \\
= & \quad \{ \text{ defn. of } \models \} \\
& \sum_{w' \in W} \{(\mathcal{A}_\pi(w, w'); \big((w' \models \rho) \cdot (w' \models \rho')\big)\} \\
= & \quad \{ (22) \} \\
& \sum_{w' \in W} \{(\mathcal{A}_\pi(w, w'); (w' \models \rho) \cdot (\mathcal{A}_\pi(w, w'); (w' \models \rho'))\}
\end{aligned}
$$

$$
\begin{aligned}
&\leq \quad \{ \text{ infimum properties} \} \\
&\quad \sum_{w' \in W} \{(\mathcal{A}_\pi(w,w'); (w' \models \rho)\} \cdot \sum_{w' \in W} \{\mathcal{A}_\pi(w,w'); (w' \models \rho'))\} \\
&= \quad \{ \text{ defn of } \models \} \\
&\quad (w \models \langle\pi\rangle\rho) \cdot (w \models \langle\pi\rangle\rho) \\
&= \quad \{ \text{ defn of } \models \} \\
&\quad (w \models \langle\pi\rangle\rho \wedge \langle\pi\rangle\rho)
\end{aligned}
$$

Therefore, $\langle\pi\rangle(\rho \wedge \rho') \rightarrow \langle\pi\rangle\rho \wedge \langle\pi\rangle\rho$.

Lemma 2. *The following are valid formulas in any $\mathcal{GDL}(\mathbf{A})$:*

(2.1) $\langle\pi + \pi'\rangle\rho \leftrightarrow \langle\pi\rangle\rho \vee \langle\pi\rangle\rho$
(2.2) $\langle\pi;\pi'\rangle\rho \leftrightarrow \langle\pi\rangle\langle\pi'\rangle\rho$
(2.3) $\langle\pi\rangle\bot \leftrightarrow \bot$

Proof. **Axiom (2.1)**

$$
\begin{aligned}
&\quad (w \models \langle\pi + \pi'\rangle\rho) \\
&= \quad \{ \text{ defn of } \models \} \\
&\quad \sum_{w' \in W} \{\mathcal{A}_{\pi+\pi'}(w,w'); (w' \models \rho)\} \\
&= \quad \{ \text{ defn of programs interpretation} \} \\
&\quad \sum_{w' \in W} \{(\mathcal{A}_\pi(w,w') + \mathcal{A}_{\pi'}(w,w')); (w' \models \rho)\} \\
&= \quad \{ (7) \} \\
&\quad \sum_{w' \in W} \{(\mathcal{A}_\pi(w,w'); (w' \models \rho) + \mathcal{A}_{\pi'}(w,w'); (w' \models \rho))\} \\
&= \quad \{ \text{ distributivity of supremum} \} \\
&\quad \sum_{w' \in W} \{(\mathcal{A}_\pi(w,w'); (w' \models \rho)\} + \sum_{w' \in W} \{\mathcal{A}_{\pi'}(w,w'); (w' \models \rho))\} \\
&= \quad \{ \text{ defn of } \models \} \\
&\quad (w \models \langle\pi\rangle\rho) + (w \models \langle\pi'\rangle\rho) \\
&= \quad \{ \text{ defn of } \models \} \\
&\quad (w \models \langle\pi\rangle\rho \vee \langle\pi'\rangle\rho)
\end{aligned}
$$

Therefore $\langle\pi + \pi'\rangle\rho \leftrightarrow \langle\pi\rangle\rho \vee \langle\pi'\rangle\rho$.

Axiom (2.2)

$$
\begin{aligned}
& (w \models \langle\pi\rangle\langle\pi'\rangle\rho) \\
= \quad & \{ \text{ defn of } \models\} \\
& \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); (w \models \langle\pi'\rangle\rho)\} \\
= \quad & \{ \text{ defn of } \models\} \\
& \sum_{w' \in W} \big\{\mathcal{A}_\pi(w, w'); \sum_{w'' \in W} \{\mathcal{A}_{\pi'}(w', w''); (w'' \models \rho)\}\big\} \\
= \quad & \{ (7)\} \\
& \sum_{w' \in W} \big\{ \sum_{w'' \in W} \{\mathcal{A}_\pi(w, w'); \mathcal{A}_{\pi'}(w', w''); (w'' \models \rho)\}\big\} \\
= \quad & \{ \text{commutativity}\} \\
& \sum_{w'' \in W} \big\{ \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); \mathcal{A}_{\pi'}(w', w''); (w'' \models \rho)\}\big\} \\
= \quad & \{ \text{since } (w'' \models \rho) \text{ is independent of } w'\} \\
& \sum_{w'' \in W} \big\{ \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); \mathcal{A}_{\pi'}(w', w'')\}; (w'' \models \rho)\big\} \\
= \quad & \{ \text{defn. of composition}\} \\
& \sum_{w'' \in W} \big\{\mathcal{A}_{\pi;\pi'}(w, w''); (w'' \models \rho)\big\} \\
= \quad & \{ \text{defn. of } \models\} \\
& (w \models \langle\pi; \pi'\rangle\rho)
\end{aligned}
$$

Therefore $\langle\pi\rangle\langle\pi'\rangle\rho \leftrightarrow \langle\pi; \pi'\rangle\rho$.

Axiom (2.3)

$$
\begin{aligned}
& (w \models \langle\pi\rangle\bot) \\
= \quad & \{ \text{defn. of } \models\} \\
& \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); (w \models \bot)\} \\
= \quad & \{ \text{defn. of satisfaction}\} \\
& \sum_{w' \in W} \{\mathcal{A}_\pi(w, w'); \bot\} \\
= \quad & \{ (9) \text{ and } \bot = 0\} \\
& \sum_{w' \in W} \{\bot\}
\end{aligned}
$$

$$
\begin{array}{ll}
= & \{\ (4)\} \\
\quad \bot &
\end{array}
$$

Therefore $\langle\pi\rangle 0 \leftrightarrow 0$.

Lemma 3. *The following are valid formulas in any $\mathcal{GDL}(\mathbf{A})$:*

(3.1) $\langle\pi\rangle\rho \rightarrow \langle\pi^*\rangle\rho$
(3.2) $\langle\pi^*\rangle\rho \leftrightarrow \langle\pi^*;\pi^*\rangle\rho$
(3.3) $\langle\pi^*\rangle\rho \leftrightarrow \langle\pi^{**}\rangle\rho$
(3.4) $\langle\pi^*\rangle\rho \leftrightarrow \rho \vee \langle\pi\rangle\langle\pi^*\rangle\rho$

Proof. **Axiom (3.1)** In order to proof this axiom we have first to observe that for any $a, b, c \in A$, $a \leq b$ implies $a;c \leq b;c$. Supposing $a \leq b$, i.e., $a + b = b$, we have that

$$a;c + b;c \ =_{\{(8)\}}\ (a + b);c \ =_{\{\text{by hypothesis } a+b=b\}}\ b;c$$

i.e., $a;c \leq b;c$. Moreover, we have also to check that $a \leq a^*$ which comes directly from (10) by monotonicity of the supremum and transitivity. Hence (and since $\mathbb{M}_n(\mathbf{A})$ is an action lattice), we have for any $w, w' \in W$,

$$
\begin{array}{ll}
& \mathcal{A}_\pi(w,w') \leq \mathcal{A}_{\pi^*}(w,w') \\
\equiv & \quad \{\ a \leq b \text{ implies } a;c \leq b;c\} \\
& \mathcal{A}_\pi(w,w');(w' \models \rho) \leq \mathcal{A}_{\pi^*}(w,w');(w' \models \rho) \\
\equiv & \quad \{\ \text{monotonicity of the supremum}\} \\
& \displaystyle\sum_{w' \in W} \{\mathcal{A}_\pi(w,w');(w' \models \rho)\} \leq \sum_{w' \in W} \{\mathcal{A}_{\pi^*}(w,w');(w' \models \rho)\} \\
\equiv & \quad \{\ \text{defn of } \models\} \\
& (w \models \langle\pi\rangle\rho) \leq (w \models \langle\pi^*\rangle\rho) \\
\equiv & \quad \{\ \text{defn of } \models\} \\
& (w \models \langle\pi\rangle\rho \rightarrow \langle\pi^*\rangle\rho)
\end{array}
$$

Therefore $\langle\pi\rangle\rho \rightarrow \langle\pi^*\rangle\rho$.

Axioms (3.2),(3.3) and (3.4) We start recalling the following well known Kleene algebra properties: $a^* = a^{**}$, $a^* = a^*;a^*$ and $1 + a;a^* = a^*$ (see [9]). By Theorem 1, we have that

$$\mathcal{A}_{\pi^*}(w,w') = \mathcal{A}_{\pi^{**}}(w,w') \tag{23}$$
$$\mathcal{A}_{\pi^*}(w,w') = \mathcal{A}_{\pi^*;\pi^*}(w,w') \tag{24}$$
$$\mathcal{A}_{1+\pi;\pi^*}(w,w') = \mathcal{A}_{\pi^*}(w,w') \tag{25}$$

The remaining of the first two proofs follows exactly the same steps of the one for Axiom **(3.1)**. For the third case,

$$
\begin{array}{ll}
& \mathcal{A}_{1+\pi;\pi^*}(w,w') = \mathcal{A}_{\pi^*}(w,w') \\
\equiv & \quad \{ \text{ program interpretation} \} \\
& \mathcal{A}_1(w,w') + \mathcal{A}_{\pi;\pi^*}(w,w') = \mathcal{A}_{\pi^*}(w,w') \\
\equiv & \quad \{ \ a \leq b \text{ implies } a;c \leq b;c \text{ in both directions} \} \\
& \big(\mathcal{A}_1(w,w') + \mathcal{A}_{\pi;\pi^*}(w,w')\big);(w' \models \rho) = \mathcal{A}_{\pi^*}(w,w');(w' \models \rho) \\
\equiv & \quad \{ \ (7) \} \\
& \mathcal{A}_1(w,w');(w' \models \rho) + \mathcal{A}_{\pi;\pi^*}(w,w');(w' \models \rho) = \mathcal{A}_{\pi^*}(w,w');(w' \models \rho) \\
\equiv & \quad \{ \text{ monotonicity of the supremum} \} \\
& \sum_{w' \in W} \{\mathcal{A}_1(w,w');(w' \models \rho) + \mathcal{A}_{\pi;\pi^*}(w,w');(w' \models \rho)\} = \sum_{w' \in W} \{\mathcal{A}_{\pi^*}(w,w');(w' \models \rho)\} \\
\equiv & \quad \{ \text{ distributivity} \} \\
& \sum_{w' \in W} \{\mathcal{A}_1(w,w');(w' \models \rho)\} + \sum_{w' \in W} \{\mathcal{A}_{\pi;\pi^*}(w,w');(w' \models \rho)\} = \sum_{w' \in W} \{\mathcal{A}_{\pi^*}(w,w');(w' \models \rho)\} \\
\equiv & \quad \{ \ \textstyle\sum_{w' \in W} \{\mathcal{A}_1(w,w');(w' \models \rho) = (w \models \rho)\} + \text{program interpretation} \} \\
& (w \models \rho) + (w \models \langle \pi;\pi^* \rangle \rho) = (w \models \langle \pi^* \rangle \rho) \\
\equiv & \quad \{ \ \textbf{(2.2)} \} \\
& (w \models \rho) + (w \models \langle \pi \rangle \langle \pi^* \rangle \rho) = (w \models \langle \pi^* \rangle \rho) \\
\equiv & \quad \{ \text{ defn of } \models \} \\
& (w \models \rho \vee \langle \pi \rangle \langle \pi^* \rangle \rho) = (w \models \langle \pi^* \rangle \rho)
\end{array}
$$

Therefore, $\langle \pi^* \rangle \rho \leftrightarrow \rho \vee \langle \pi \rangle \langle \pi^* \rangle \rho$.

4 Conclusions

The method introduced in this paper is able to generate several dynamic logics useful for the working Software Engineer. Some of them are documented in the literature, others freshly new. For instance, for verification of imperative programs, we may consider a logic whose states are valuations of program variables. Hence, and as usual, atomic programs become assignments of variables. In this context, a transition $w \rightarrow^{x:=a} w'$ means that the state w' differs from w just in the value of variable x, i.e., that $w'(x) = a$ and for any variable $y \neq x$, $w(y) = w'(y)$.

A very natural direction for future work is to enrich this framework with tests, i.e., programs $?cond$ interpreted as $\mathcal{A}_{?cond} = \{(w,w) | w \models cond\}$. As usual, this provides a way to express *if-then-else* statements in dynamic logics. Another topic deserving attention is the characterisation of program refinement

in this setting, witnessed by some class of action lattice morphisms.

Acknowledgements.

This work is financed by the ERDF - European Regional Development Fund through the COMPETE Programme (operational programme for competitiveness) and by National Funds through the FCT - Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) within projects FCOMP-01-0124-FEDER-037281 and FCOMP-01-0124-FEDER-028923 and by project NORTE-07- 0124-FEDER-000060, co-financed by the North Portugal Regional Operational Programme (ON.2), under the National Strategic Reference Framework (NSRF), through the European Regional Development Fund (ERDF).

References

1. B. Beckert. A dynamic logic for the formal verification of java card programs. In I. Attali and T. P. Jensen, editors, *Java Card Workshop*, volume 2041 of *Lecture Notes in Computer Science*, pages 6–24. Springer, 2000.
2. M. Droste and P. Gastin. Weighted automata and weighted logics. *Theor. Comput. Sci.*, 380(1-2):69–86, 2007.
3. H. Furusawa. The categories of kleene algebras, action algebras and action lattices are related by adjunctions. In R. Berghammer, B. M oller, and G. Struth, editors, *RelMiCS*, volume 3051 of *Lecture Notes in Computer Science*, pages 124–136. Springer, 2003.
4. S. F. Goble. Grades of modality. *Logique et Analyse*, 13:323–334, 1970.
5. S. Gottwald. *A Treatise on Many-Valued Logics*. Studies in Logic and Computation (volume 9). Research Studies Press, 2001.
6. D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. MIT Press, 2000.
7. D. Kozen. On action algebras. manuscript in: Logic and Flow of Information, Amsterdam, 1991.
8. D. Kozen. A probabilistic PDL. *J. Comput. Syst. Sci.*, 30(2):162–178, 1985.
9. D. Kozen. A completeness theorem for kleene algebras and the algebra of regular events. *Inf. Comput.*, 110(2):366–390, 1994.
10. B. Lopes, M. Benevides, and E. H. Haeusler. Propositional dynamic logic for petri nets. *Logic Journal of IGPL*, 2014.
11. O. Mürk, D. Larsson, and R. Hähnle. Key-c: A tool for verification of c programs. In F. Pfenning, editor, *CADE*, volume 4603 of *Lecture Notes in Computer Science*, pages 385–390. Springer, 2007.
12. N. J. Nilsson. Probabilistic logic. *Artif. Intell.*, 28(1):71–87, 1986.
13. A. Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, 2010.
14. A. Platzer. A complete axiomatization of quantified differential dynamic logic for distributed hybrid systems. *Logical Methods in Computer Science*, 8(4), 2012.
15. V. R. Pratt. Action logic and pure induction. In *JELIA*, volume 478 of *Lecture Notes in Computer Science*, pages 97–120. Springer, 1990.
16. S. Suh, U. Tanik, Carbone, and A. J.N., Eroglu. *Applied Cyber-Physical Systems*. Springer Verlag, 2014.
17. W. van der Hoek. On the semantics of graded modalities. *Journal of Applied Non-Classical Logics*, 2(1), 1992.

Model-Driven Engineering in the Heterogeneous Tool Set

Daniel Calegari[1], Till Mossakowski[2], and Nora Szasz[3]

[1] Universidad de la República, Uruguay
dcalegar@fing.edu.uy
[2] Otto-von-Guericke University Magdeburg, Germany
mossakow@iws.cs.uni-magdeburg.de
[3] Facultad de Ingeniería, Universidad ORT Uruguay
szasz@ort.edu.uy

Abstract. We have defined a unified environment that allows formal verification within the Model-Driven Engineering (MDE) paradigm using heterogeneous verification approaches. The environment is based on the Theory of Institutions, which provides a sound basis for representing MDE elements and a way for specifying translations from these elements to other logical domains used for verification, such that formal experts can choose the domain in which they are more skilled to address a formal proof. In this paper we present how this environment can be supported in practice by the Heterogeneous Tool Set (Hets). We define semantic-preserving translations from the MDE elements to the core language of Hets, and we also show how it is possible to move from it to other logics, both to supplement the original specification with other verification properties and to perform a heterogeneous verification.

Keywords: verification, formal methods, Model-Driven Engineering

1 Introduction

The Model-Driven Engineering (MDE,[1]) paradigm is based on the construction of models representing different views of the system to be constructed, and model transformations as the main activity within the software development process. In this context, there are multiple properties that can be verified [2], from syntactic to semantic ones, and at different abstraction levels. Whenever formal verification is mandatory, there is a plethora of verification approaches with different objectives, formalisms and supporting tools, which are heterogeneous and not integrated. With an heterogeneous approach [3], different formalisms are used for expressing parts of a problem and semantic-preserving mappings allow the communication between these formalisms in order to compose different views to an overall specification of the whole problem. We have followed this approach by proposing a theoretical environment for the formal verification of different MDE aspects using heterogeneous verification approaches [4], based on the theory of Institutions [5]. This environment proposes a generic representation of the MDE

elements (by means of institutions) which can be formally (and automatically) translated into other formalisms, providing the "glue" that formal experts need to choose the formalism in which they are more skilled to address a formal proof.

In this paper we show how the environment can be supported in practice using the Heterogenous Tool Set (Hets,[3,6]), which is meant to support heterogeneous multi-logic specifications. It also provides proof management capabilities for monitoring the overall correctness of a heterogeneous specification whereas different parts of it are verified using (possibly) different formalisms. We first define from a theoretical perspective how MDE elements can be integrated in this tool by defining semantic-preserving translations to the Common Algebraic Specification Language (Casl,[7]), which is the core language of Hets. The existent connections between Casl and other formalisms broadens the spectrum of formal domains in which verification can be addressed. We also detail the implementation of a prototype which allows us to specify MDE elements, supplement them with multi-logic properties, and perform a heterogeneous verification.

The remainder of the paper is structured as follows. In Section 2 we introduce the main concepts of MDE based on a running example, and in Section 3 we summarize how these elements can be represented within our institution-based environment. Then, in Section 4 we present how this environment can be formally connected with Casl, and in Section 5 we give details about an implementation of these ideas using Hets. Finally, in Section 6 we discuss related work and in Section 7 we present some conclusions and an outline of further work.

2 Model-Driven Engineering

In MDE there are two key elements: models specifying different views of the system to be constructed and model transformations allowing the (semi)automatic construction of the system by processing the models.

Every model *conforms* to a metamodel which introduces the syntax and semantics of certain kinds of models. The MetaObject Facility (MOF, [8]) is a standard language for metamodeling, basically defining hierarchical-structured classes with properties that can be attributes (named elements with an associated primitive type or class) or associations (relations between classes in which each class plays a role within the relation). Every property has a multiplicity which constraints the number of elements that can be related through it. If there are conditions that cannot be captured by the structural rules of this language, the Object Constraint Language (OCL, [9]) is used to specify them. These considerations allow defining conformance in terms of *structural* and *non-structural* conformance. Structural conformance with respect to a metamodel means that in a given model every object and link is well-typed and the model also respects the multiplicity constraints. Non-structural conformance means that a given model respects the invariants specified with the supplementary language.

Consider as an example a simplified version of the well-known Class to Relational model transformation [10]. The metamodel in the left side of Figure 1 defines UML class diagrams, where classifiers (classes and primitive types) are

contained in packages. Classes can contain one or more attributes and may be declared as persistent, and each attribute is typed by a primitive type. Notice that a class must contain only one or two attributes, and also that the Classifier class is not abstract. We handle these aspects differently from UML class diagrams in order to have a more complete example. In the right side of Figure 1 there is a model composed by a persistent class of name `ID` within a package of name `Package`. The class has an attribute of name `value` and type `String`.

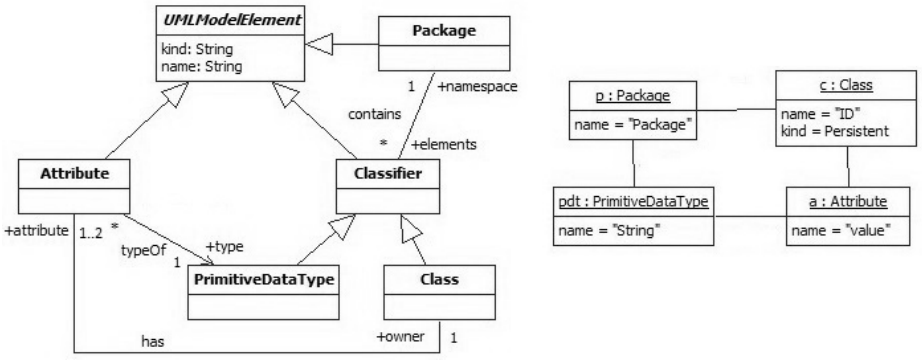

Fig. 1. Class metamodel and model of the example

A model transformation takes as input a model conforming to certain metamodel and produces as output another model conforming to another metamodel (possibly the same). Query/View/Transformation Relations (QVT-Relations, [10]) is a relational language which defines transformation rules as mathematical relations between source and target elements. A transformation is a set of interconnected relations: top-level relations that must hold in any transformation execution, and non-top-level relations that are required to hold only when they are referred from another relation. Every relation defines a set of variables, and source and target patterns which are used to find matching sub-graphs of elements in a model. Relations can also contain a `when` clause which specifies the conditions under which the relationship needs to hold, and a `where` clause which specifies the condition that must be satisfied by all model elements participating in the relation. The `when` and `where` clauses, as well as the patterns may contain arbitrary boolean OCL expressions and can invoke other relations.

The transformation of the example basically describes how persistent classes within a package are transformed into tables within a schema, and attributes of a class are transformed into columns of the corresponding table. Below we show an excerpt of this transformation. There are keys defined as the combination of those properties of a class that together can uniquely identify an instance of that class, e.g. there are no two tables with the same name within the same schema.

```
transformation uml2rdbms ( uml : UML , rdbms : RDBMS ) {
  key RDBMS::Table {name, schema};

  top relation PackageToSchema { ... }

  top relation ClassToTable {
    cn, prefix : String;
    checkonly domain uml c : UML::Class {
      namespace = p : UML::Package {}, kind = 'Persistent', name = cn
    };
    enforce domain rdbms t : RDBMS::Table {
      schema = s : RDBMS::Schema {}, name = cn
    };
    when { PackageToSchema(p, s); }
    where { AttributeToColumn(c, t);}
  }

  relation AttributeToColumn { ... }
}
```

3 An Institution-Based Environment for MDE

Our environment [4] is based on representing models (from now on SW-models to avoid confusion), metamodels, the conformance relation, transformations and verification properties in some consistent and interdependent way without depending on any specific logical domain. We follow an heterogeneous specification approach [3] which is based on providing *Institutions* [5] for representing the syntax and semantics of the elements. An institution is defined as:

- a category Sign of signatures (vocabularies for constructing sentences in a logical system) and signature morphisms (translations between vocabularies)
- a functor Sen : Sign $\rightarrow$ Set giving a set of sentences for each signature and a function Sen(σ):Sen(Σ_1)$\rightarrow$ Sen(Σ_2) translating formulas to formulas for each signature morphism $\sigma : \Sigma_1 \rightarrow \Sigma_2$;
- a functor Mod : Sign$^{op} \rightarrow$ **Cat** , giving a category Mod(Σ) of models (providing semantics) for each signature Σ and a reduct functor Mod(σ):Mod(Σ_2)$\rightarrow$ Mod(Σ_1) translating models to models (and morphisms to morphisms) for each signature morphism;
- a satisfaction relation of sentences by models, such that when signatures are changed (by a signature morphism), satisfaction of sentences by models changes consistently, i.e. $M_2 \models_{\Sigma_2}$ Sen(σ)(φ) iff Mod(σ)(M_2) $\models_{\Sigma_1} \varphi$

We provide an institution $\mathcal{I}^{\mathrm{Q}}$ for QVT-Relations check-only unidirectional transformations (which we called QVTR). This institution needs a representation of SW-models and metamodels, therefore we define an institution $\mathcal{I}^{\mathrm{M}}$ for the structural conformance relation between them based on a simplified version of MOF (which we called CSMOF). Complete definitions can be found in [11].

The institution $\mathcal{I}^{\mathrm{M}}$ represents the MOF-based structural conformance relation between metamodels and SW-models. From any metamodel we can derive a signature $\Sigma = (\boldsymbol{C}, \alpha, \boldsymbol{P})$ declaring: a finite class hierarchy $\boldsymbol{C} = (C, \leq_C)$ (a partial order between classes representing the inheritance relation between them) extended with a subset $\alpha \subseteq C$ denoting abstract classes; and a properties declaration (attributes and associations) $\boldsymbol{P} = (R, P)$ where R is a finite set of role names with a default role name "_", and P is a finite set of properties of the form $\langle r_1 : c_1, r_2 : c_2 \rangle$ representing a property and its opposite. The type c_i attached to the role r_i represents the type of the property, as well the type in the opposite side represents its owned class. By $\boldsymbol{T}(\boldsymbol{C})$ we denote the *type extension* of $\boldsymbol{C}$ by primitive types (e.g. Boolean) and type constructors (e.g. List). Formulas represent multiplicity constraints determining whether the number of elements in a property end is bounded (upper and/or lower). They are defined as follows: $\Phi ::= \#C \bullet R = \mathbb{N} \mid \mathbb{N} \leq \#C \bullet R \mid \#C \bullet R \leq \mathbb{N}$ The #-expressions return the number of links in a property when some role is fixed. The $\bullet$ operator represents the selection of the elements linked with another of class C through a role in R. An interpretation $\mathcal{I}$ (or model) contains a semantic representation for a SW-model, i.e. objects and links. It consists of a tuple $(\boldsymbol{V}^{\boldsymbol{T}}_{\boldsymbol{C}}(\boldsymbol{O}), \boldsymbol{A})$ where $\boldsymbol{V}^{\boldsymbol{T}}_{\boldsymbol{C}}(\boldsymbol{O}) = (V_c)_{c \in T(C)}$ is a $\boldsymbol{T}(\boldsymbol{C})$-object domain (a family of sets of object identifiers), $\boldsymbol{A}$ contains a relation $\langle r_1 : c_1, r_2 : c_2 \rangle^{\mathcal{I}} \subseteq V_{c_1} \times V_{c_2}$ for each relation name $\langle r_1 : c_1, r_2 : c_2 \rangle \in P$ with $c_1, c_2 \in T(C)$, and $c_2 \in \alpha$ implies $O_{c_2} = \bigcup_{c_1 \leq_C c_2} O_{c_1}$. Finally, an interpretation $\mathcal{I}$ satisfies a formula φ with some $c \bullet r$ if for any object of class c, the number of elements within $\mathcal{I}$ related through the role r (of a property of the class c) satisfies the multiplicity constraints. The satisfaction relation checks the multiplicity requirements of the structural conformance relation.

The institution $\mathcal{I}^{\mathrm{Q}}$ represents QVT-Relations transformations by extending the CSMOF institution. A signature is a pair $\langle \Sigma^{\mathrm{M}}_1, \Sigma^{\mathrm{M}}_2 \rangle$ representing the source and target metamodels of the transformation, and an interpretation is a tuple $\langle \mathcal{M}^{\mathrm{M}}_1, \mathcal{M}^{\mathrm{M}}_2 \rangle$ of disjoint $\mathrm{Sign}^{\mathrm{M}}_i$-interpretations that contains a semantic representation for the source and target SW-models. A formula φ^{K} represents a key constraint of the form $\langle c, \{r_1, ..., r_n\} \rangle$ $(1 \leq n)$ with $c \in C_i$ (i = 1..n) a class in one of the metamodels, $r_j \in R_i$ (j = 1..n) roles defined in properties in which such class participates (having such role or at the opposite side of it). Roles determine the elements within these properties that together can uniquely identify an instance of the class. A formula φ^{R} represents a set of interrelated transformation rules, such that, given variables $X^s = (X^s)_{s \in (\bigcup_i T(C_i))}$, the formula is a finite set of tuples representing rules of the form $\langle \mathrm{top}, \mathrm{VarSet}, \mathrm{ParSet}, \mathrm{Pattern}_1, \mathrm{Pattern}_2, \mathrm{when}, \mathrm{where} \rangle$, where top $\in \{true, false\}$ defines if the rule is a top-level relation or not, VarSet $\subseteq X^s$ is the set of variables used within the rule, ParSet $\subseteq$ VarSet representing the set of variables taken as parameters when the rule is called from another one, $\mathrm{Pattern}_i$ $(i = 1, 2)$ are the source and target patterns, and when/where are the `when`/`where` clauses of the rule, respectively. A pattern is a tuple $\langle E_i, A_i, Pr_i \rangle$ such that $E_i \subseteq (X^c)_{c \in C_i}$ is a set of class-indexed variables, A_i is a set of elements representing associations of the form $rel(p, x, y)$ with $p \in P_i$ and $x, y \in E_i$,

and Pr_i is a predicate over these elements. A `when`/`where` clause is a pair $\langle \text{when}_c, \text{when}_r \rangle$ such that when_c is a predicate with variables in VarSet, and when_r is a set of pairs of transformation rules and their parameters. The satisfaction relation expresses that the target SW-model is the result of transforming the source SW-model (both within the interpretation) according to the transformation rules and also that key constraints hold (both represented as formulas).

Institutions can be formally connected by means of (co)morphisms. Then, by defining these semantic-preserving translations, it is possible to connect MDE elements to potentially several logics for formal verification. In this way, we just specify MDE elements once, then spread this information into other logics to supplement this specification with additional properties, and finally choose the verification approach we want to use. To the extent that there are many logics connected through comorphisms, the capabilities of our environment increases. The environment supports a separation of duties between software developers (MDE and formal methods experts) such that a formal perspective is available whenever it is required. Moreover, comorphisms can be automated, as we show in the following sections, thus the environment is scalable in terms of the rewriting of MDE elements in each logic. Although our proposal is aligned with OMG standards, this idea can be potentially formalized for any transformation approach and language. This allows extending the approach as far as necessary.

4 Borrowing Proof Capabilities

We make use of the possibility of connecting our institutions to potentially several host logics, each one with its own proof system. The host logic allows both to supplement the information contained within the MDE elements with properties specified in the host logic, and to borrow its proof calculus for formal verification. For this, we us *generalized theoroidal comorphisms* (GTC,[12]). A GTC between two institutions $\mathcal{I}$ and $\mathcal{J}$consists of a functor $\Phi : \text{Th}^{\mathcal{I}} \rightarrow \text{Th}^{\mathcal{J}}$ translating theories (pairs of signatures and set of sentences), and a natural transformation $\beta : (\Phi)^{op}; \text{Mod}^{\mathcal{J}} \rightarrow \text{Mod}^{\mathcal{I}}$ translating models in the opposite direction.

We do not define GTC from the institutions defined in the last section, but from extended institutions $\mathcal{I}^{\text{M}^+}$ and $\mathcal{I}^{\text{Q}^+}$. We extend the definition of CSMOF formulas with a syntactic representation of SW-models as follows:

$$\Omega ::= x^c \mid \langle r_1, {x_1}^{c_1}, r_2, {x_2}^{c_2} \rangle \mid \Omega \oplus \Omega$$

with $x^c \in X^c$ a variable representing a typed element, $\langle r_1, {x_1}^{c_1}, r_2, {x_2}^{c_2} \rangle$ representing a link between two typed elements with their respective roles, and $\Omega \oplus \Omega$ the composition of these elements. In the case of QVTR, we extend QVTR formulas by including extended CSMOF formulas, i.e. now there is a representation of multiplicity constraints and SW-models, indexed by the institutions in which they are defined. These extensions make it possible to use a proof system such that it is possible to prove that constraints (as a formula) are derived from a syntactic representation of a SW-model, which is the context where the verification must be done. An exhaustive discussion on this topic can be found in [11].

We defined GTCs from our extended institutions to CASL, a general-purpose specification language. The institution $\mathcal{I}^{\mathrm{C}}$ underlying CASL is the sub-sorted partial first-order logic with equality and constraints on sets $SubPCFOL^{=}$, a combination of first-order logic and induction with subsorts and partial functions. Since CASL has a sound proof calculus, and our comorphisms admit borrowing of entailment [3], we can translate our proof goals using the comorphism into CASL and use its proof calculus also for proving properties of our extended CSMOF and QVTR specifications. The importance of CASL is that it is the main language within the Heterogenous Tool Set (HETS, [3]), a tool meant to support heterogeneous multi-logic specifications. This comorphism not only allows us to have tool support for the verification of model transformation by using HETS (as will be introduced in Section 5) but also to move between the graph of logics within HETS to take advantage of the benefits of each logic.

In what follows we introduce CASL and resume the encoding of the main components of the extended institutions into it. An example of the encoding is given in Section 5.1, and a complete version can be found in [11].

4.1 Common Algebraic Specification Language

The institution $\mathcal{I}^{\mathrm{C}}$ for CASL is defined as follows. Signatures consist of a set S of sorts with a subsort relation $\leq$ between them, together with a family $\{PF_{w,s}\}_{w\in S^*,s\in S}$ of partial functions, $\{TF_{w,s}\}_{w\in S^*,s\in S}$ of total functions and $\{P_w\}_{w\in S^*}$ of predicate symbols. Signature morphisms consist of maps taking sort, function and predicate symbols respectively to a symbol of the same kind in the target signature, and they must preserve subsorting, typing of function and predicate symbols and totality of function symbols.

For a signature Σ, terms are formed starting with variables from a sorted set X using applications of function symbols to terms of appropriate sorts, while sentences are partial first-order formulas extended with *sort generation constraints* which are triples (S', F', σ') such that $\sigma' : \Sigma' \to \Sigma$ and S' and F' are respectively sort and function symbols of Σ'. Models interpret sorts as non-empty sets such that subsorts are injected into supersorts, partial/total function symbols as partial/total functions and predicate symbols as relations.

The satisfaction relation is the expected one for partial first-order sentences. A sort generation constraint (S', F', σ') holds in a model M if the carriers of the reduct of M along σ' of the sorts in S' are generated by function symbols in F'.

4.2 Encoding CSMOF into CASL

We define a GTC between the extended CSMOF institution $\mathcal{I}^{\mathrm{M}^+}$ and the institution $\mathcal{I}^{\mathrm{C}}$ for $SubPCFOL^{=}$. The class hierarchy represented within a $\mathcal{I}^{\mathrm{M}^+}$ signature is basically translated into a set of sorts complying with a subsorting relation, properties are translated into predicates, and an axiom is introduced to relate predicates derived from bidirectional properties. Formally, every $\mathcal{I}^{\mathrm{M}^+}$ signature $\Sigma = (\boldsymbol{C}, \alpha, \boldsymbol{P})$ with $\boldsymbol{C} = (C, \leq_C)$ and $\boldsymbol{P} = (R, P)$ is translated into a theory $((S, TF, PF, P, \leq_S), E)$ such that:

- For every class name c in C, there is a sort name $c \in S$.
- For every $c_1 \leq_C c_2$ with $c_1, c_2 \in C$, we have $c_1 \leq_S c_2$ with $c_1, c_2 \in S$.
- For every $c \in \alpha$ there is an axiom in E stating that c is the disjoint embedding of its subsorts (sort generation constraint).
- For every $\langle r_1 : c_1, r_2 : c_2 \rangle \in P$, there are two predicates $r_1 : c_2 \times c_1$ and $r_2 : c_1 \times c_2 \in \Pi$, and an axiom in E stating the equivalence of the predicates, i.e. $r_1(x, y)$ iff $r_2(y, x)$ with $x \in S_1, y \in S_2$. In the case of predicates with the default role name _, we only generate the predicate in the opposite direction of the default role, i.e. if $\langle \_ : c_1, r_2 : c_2 \rangle$ or $\langle r_1 : c_1, \_ : c_2 \rangle$ we only have $r_2 : c_1 \times c_2$ or $r_1 : c_2 \times c_1$, respectively.

We consider the existence of a built-in extension of the institution $\mathcal{I}^{\mathrm{C}}$, e.g. the CASL standard library. The sets of functions TF and PF within this extension contain those functions defined for built-in types (like + for strings).

As an example, the signature corresponding to the class metamodel in Figure 1 is translated into a theory such that there are sorts for each class, e.g. UMLModelElement and Package, within the subsorting relation, e.g. Package $\leq_S$ UMLModelElement; and there are predicates for each property, e.g. elements : Package $\times$ Classifier and name : UMLModelElement $\times$ String. There is a sort generation constraint stating that UMLModelElement is the disjoint embedding of its subsorts Attribute, Classifier, and Package. There are also axioms stating the equivalence of the predicates derived from bidirectional properties, e.g. $\forall\, x : \text{Package}, y : \text{Classifier}.\ \text{elements}(x, y) \Leftrightarrow \text{namespace}(y, x)$

In the case of a SW-model formula Ω, each variable within the formula (representing an object) is translated into a total function of the corresponding type. We also add several axioms in order to represent implicit constraints in the $\mathcal{I}^{\mathrm{M}^+}$ institution which are not necessarily kept when representing the basic elements in $SubPCFOL^{=}$, as for example the need of distinguishing between two different variables (functions in the target institution) and the specification of the cases in which a property holds (when there is a syntactic link represented within the formula Ω). Formally,

- For every $x^c \in \upsilon(\Omega)$ there is a total function $x : c \in TF$ with $c \in S$
- For every $\langle r_1, x^{c_1}, r_2, y^{c_2} \rangle \in \omega(\Omega)$ with $\langle r_1 : c_1, r_2 : c_2 \rangle \in P$, there is an axiom in E stating that the predicate $r_2 : c_1 \times c_2$ holds for $x : c_1, y : c_2 \in TF$. Notice that the opposite direction holds by the equivalence of predicates stated during the signature translation.
- E has some additional axioms:
  - Distinguishability: $\{x_i \neq x_j \mid i \neq j.\ x_i, x_j : c \in TF\}$ for all $c \in S$
  - Completeness of elements: for all $x : c$ we have that $x = o_i$ for some $o_i : c \in TF$. When c is a non-abstract class having sub-classes, completeness must be defined for $o_i : c' \in TF$ for all $c' \leq c$.
  - Completeness of relations: for all $x : c_1, y : c_2$ we have that $r(x, y)$ holds only if $x = o_1$ and $y = o_2$ for some $o_1 : c_1, o_2 : c_2$ for which $r(c_1, c_2)$ hold.

The "distinguishability" and "completeness of elements" axioms correspond to the "no junk, no confusion" principle: there are no other values than those denoted by the functions $x : c$, and distinct functions denote different values.

The variables within the class SW-model in Figure 1 are translated into total functions, e.g. p : Package, c : Class and ID : String. Moreover, for every link there is an axiom stating that the corresponding predicate holds for the functions corresponding to the translated elements within the link. This axiom can be stated in conjunction with the "completeness of relations", e.g. $\forall\, x : \text{Package}, y : \text{Classifier}.\ \text{elements}(x, y) \Leftrightarrow (x = \text{p} \ \wedge\ y = \text{c}) \vee (x = \text{p} \ \wedge\ y = \text{pdt})$. In the case of the non-abstract class Classifier which has sub-classes, the "completeness of elements" constraint is stated by the axiom: $\forall\, x : \text{Classifier}.\ x = \text{c} \vee x = \text{pdt}$. Finally, the "distinguishability" constraint must be stated between elements of sorts related by the subsorting relation. For example, in the case of the elements within the UMLModelElement hierarchy, we have the following constraint:
$\neg(\text{a} = \text{c}) \wedge \neg(\text{a} = \text{p}) \wedge \neg(\text{a} = \text{pdt}) \wedge \neg(\text{c} = \text{p}) \wedge \neg(\text{c} = \text{pdt}) \wedge \neg(\text{p} = \text{pdt})$.

For the translation of a multiplicity constraint formula we define the following predicates for constraining the size of the set of elements in a relation:

- $min(n, R : D \times C)$ holds if for all $y : D$ there exists $x_1, ..., x_n : C$ such that $R(y, x_i)$ for all $i = \{1..n\}$, and $x_i \neq x_j$ for all $i = \{1..n-1\}$, $j = i + 1$.
- $max(n, R : D \times C)$ holds if for all $y : D$ and $x_1, ..., x_{n+1} : C$, $Rel(y, x_i)$ for all $i = \{1..n+1\}$ implies there is some $x_i = x_j$ with $i = \{1..n\}$, $j = i + 1$.

The first predicate states that there are at least n different elements related to every element y by the relation R, which represents a minimal cardinality for the relation. The other predicate states that there are no more than n elements related to any element y by the relation R, which represents a maximal cardinality for the relation. Using these predicates, we can translate any multiplicity constraint formula as follows:

- $n \leq \#D \bullet R$ is translated into $min(n, R : D \times C)$
- $\#D \bullet R \leq n$ is translated into $max(n, R : D \times C)$
- $\#D \bullet R = n$ is translated into $min(n, R : D \times C) \wedge max(n, R : D \times C)$

such that $Q : C \times D, R : D \times C \in \Pi$ are the predicates generated by the translation of the property $\langle R : C, Q : D \rangle$. If the multiplicity constraint involves the other end, i.e. $C \bullet Q$, the predicate $Q : C \times D$ is used instead of $R : D \times C$.

As an example, the formula #(UMLModelElement $\bullet$ name) = 1 derived from Figure 1 is translated into the conjunction of

$$
\begin{aligned}
&min(1, \text{name} : \text{UMLModelElement} \times \text{String}) = \\
&\qquad \forall\, x_1 : \text{UMLModelElement}.\ \exists\, y_1 : \text{String}.\ \text{name}(x_1, y_1) \\
&max(1, \text{name} : \text{UMLModelElement} \times \text{String}) = \\
&\qquad \forall\, x_1 : \text{UMLModelElement}, y_2, y_1 : \text{String}. \\
&\qquad\qquad (\text{name}(x_1, y_1) \wedge\ \text{name}(x_1, y_2)) \Rightarrow y_1 = y_2
\end{aligned}
$$

Given a $\mathcal{I}^{\text{M}^+}$ theory $T = \langle \Sigma, \Psi \rangle$, a $\mathcal{I}^{\text{C}}$ model M of its translated theory (Σ', E) is translated into a Σ-interpretation denoted $I = (\boldsymbol{V}_{\boldsymbol{C}}^{\boldsymbol{T}}(\boldsymbol{O}), \boldsymbol{A})$ such that: each non-empty carrier set $|M|_s$ with $s \in S$, is translated into the set V_c in the object domain $\boldsymbol{V}_{\boldsymbol{C}}^{\boldsymbol{T}}(\boldsymbol{O})$, with s the translation of type $c \in T(C)$; and each relation p_M of a predicate symbol $r_2(c_1, c_2) \in P$ derived from the translation of a predicate $\langle r_1 : c_1, r_2 : c_2 \rangle$, is translated into the relation $p^{\mathcal{I}} \subseteq V_{c_1} \times V_{c_2} \ \in \boldsymbol{A}$.

4.3 Encoding QVTR into CASL

We define a GTC between the extended QVTR institution $\mathcal{I}^{\mathrm{Q}^+}$ and the institution $\mathcal{I}^{\mathrm{C}}$ for $SubPCFOL^{=}$. Every $\mathcal{I}^{\mathrm{Q}^+}$ signature $\langle \Sigma_1^{\mathrm{M}}, \Sigma_2^{\mathrm{M}} \rangle$ is translated by the functor Φ into a theory such that each signature Σ_i^{M} is translated as defined in the encoding of CSMOF into CASL. We assume that the institution $\mathcal{I}^{\mathrm{E}}$ of the expressions language has a correspondence (via a comorphism) with the built-in extension of the institution $\mathcal{I}^{\mathrm{C}}$.

Formulas representing keys and transformation rules are translated into named first-order formulas. Formulas will be of the form $P \Leftrightarrow F$ such that P is the predicate naming the formula, and F represents the conditions which must hold in order to satisfy a key constraint φ^{K} or transformation φ^{R}.

In the case of a formula φ^{K}, the formula F defines that there are not two different instances of that class with the same combination of properties conforming the key of such class. Formally, any formula $\langle C, \{r_1, ..., r_n\} \rangle$ is translated into a predicate key_C naming a key constraint definition, and a formula of the form $key\_C \Leftrightarrow \forall x, y \ \in C, v_j : T_j.\ x \neq y \rightarrow \bigwedge_{i,j} r_i(x, v_j) \ \rightarrow \bigvee_{i,j} \neg r_i(y, v_j)$, with $r_i(\_, \_)$ one of the two predicates from the translation of the property $\langle r_1 : C_1, r_2 : C_2 \rangle$ such that one of the roles is of type C and the other of type T_j.

The key formula in the example is translated into the expression

$$\begin{aligned} &\text{key\_Table} \Leftrightarrow \forall x, y \ \in \text{Table}, v_1 : \text{String}, \ v_2 : \text{Schema}. \\ &\quad x \neq y \rightarrow \text{name}(x, v_1) \wedge \text{schema}(x, v_2) \rightarrow \neg\text{name}(y, v_1) \vee \neg\text{schema}(y, v_2) \end{aligned}$$

In the case of a formula φ^{R}, the formula F declares that top-level relations must hold, and each individual rule is translated into the set of conditions stated by the checking semantics of QVT-Relations. Formally, every rule $\text{Rule} = \langle \text{top}, \text{VarSet}, \text{ParSet}, \text{Pattern}_i \ (i = 1, 2), \text{when}, \text{where} \rangle \in \varphi^{\mathrm{R}}$ is translated into: a predicate $Rule : T_1 \times ... \times T_n \in P$ with $\text{ParSet} = \{T_1, .., T_n\}$, and a predicate Top_Rule without parameters (only if $\text{top} = true$), naming the formula; and a formula $\forall v_1 : T_1, ..., v_n : T_n.\ Rule(v_1, ..., v_n) \Leftrightarrow F$ such that $Rule(v_1, ..., v_n)$ is the predicate defined before. In the case of a top rule, there is also a formula $Rule \Leftrightarrow F$. For the formula F there are two cases corresponding to the checking semantics of QVT-Relations:

1. If $\text{WhenVarSet} = \emptyset$

$$\begin{aligned} &\forall\ x_1, ..., x_n \in \ (\text{VarSet} \backslash \text{2\_VarSet}) \backslash \text{ParSet}.\ (\Phi(\text{Pattern}_1) \rightarrow \\ &\quad \exists\ y_1, ..., y_m \in \text{2\_VarSet} \backslash \text{ParSet}.\ (\Phi(\text{Pattern}_2) \ \wedge \ \Phi(\text{where}))) \end{aligned}$$

2. If $\text{WhenVarSet} \neq \emptyset$

$$\begin{aligned} &\forall\ z_1, ..., z_o \in \ \text{WhenVarSet} \backslash \text{ParSet}.\ (\Phi(\text{when}) \rightarrow \\ &\quad \forall\ x_1, ..., x_n \in \ (\text{VarSet} \backslash (\text{WhenVarSet} \cup \text{2\_VarSet})) \backslash \text{ParSet}. \\ &\qquad (\Phi(\text{Pattern}_1) \rightarrow \exists\ y_1, ..., y_m \in \text{2\_VarSet} \backslash \text{ParSet}. \\ &\qquad\qquad (\Phi(\text{Pattern}_2) \ \wedge \ \Phi(\text{where})))) \end{aligned}$$

The translation of $\text{Pattern}_i = \langle E_i, A_i, Pr_i \rangle$ is the formula $\bigwedge r_2(x,y) \wedge \Phi(Pr_i)$ such that $r_2(x,y)$ is the translation of predicate $p = \langle r_1 : C, r_2 : D \rangle$ for every $rel(p,x,y) \in A_i$ with $x : C, y : D$; and $\Phi(Pr_i)$ is the translation of the predicate into CASL. Moreover, the translation of $\text{when} = \langle \text{when}_c, \text{when}_r \rangle$ (or where) is the formula $\bigwedge Rule(v) \wedge \Phi(\text{when}_c)$ such that $Rule(v)$ is the parametric invocation of the rule $(\text{Rule}, v) \in \text{when}_r$, and $\Phi(\text{when}_c)$ is the translation of the predicate.

Back to the example, for each rule there is a predicate defining the rule. The relation `ClassToTable` is translated into the expression (in CASL syntax):

```
Top_ClassToTable <=> forall p : Package; s : Schema
                     . PackageToSchema(p,s) =>
                       forall c : Class; cn : String . namespace(c,p)
                         /\ kind(c,Persistent) /\ name(c,cn) =>
                           exists t : Table . schema(t,s)
                             /\ name(t,cn) /\ AttributeToColumn(c,t)
```

This formula says that the top-level relation holds whereas for every package and schema satisfying the relation `PackageToSchema`, if there is a persistent class within that package, there must exists a table in the corresponding schema with the same class name. Moreover, the attributes and columns of both elements must be in the relation `AttributeToColumn`.

Given a $\mathcal{I}^{\mathrm{Q}^+}$ theory $T = \langle \Sigma, \Psi \rangle$, a model M of its translated theory (Σ', E) is translated into a Σ-model $\mathcal{M} = \langle \mathcal{M}_1^{\mathrm{M}}, \mathcal{M}_2^{\mathrm{M}} \rangle$ by constructing disjoint models with an interpretation of elements for each corresponding $\mathcal{I}^{\mathrm{M}^+}$ theory. Each $\mathcal{M}_i^{\mathrm{M}}$ $(i = 1, 2)$ is defined as in Section 4.2.

5 The Environment in Action

We have implemented a prototype of our environment using the Heterogeneous Tools Set (HETS,[3,6]). HETS is an open source software providing a general framework for formal methods integration and proof management, based on the Theory of Institutions, as introduced above. Based on this foundation, HETS supports a variety of different logics. More specifically, HETS consists of logic-specific tools for the parsing and static analysis of basic logical theories written in the different involved logics (e.g. our extended CSMOF and QVTR institutions), as well as a logic-independent parsing and static analysis tool for structured theories and theory relations. Proof support for other logics can be obtained by using logic translations defined by comorphisms (e.g. from CSMOF to CASL). Our prototype and examples can be downloaded together with the HETS distribution.

Within this prototype, MDE experts can specify model transformations in their domain and such specifications can be complemented by verification experts with other properties to be verified, e.g. non-structural constraints. All this information is taken by HETS, which performs automatic translations of proof obligations into other logics and allows selecting the corresponding prover to be used, whilst a graphical user interface is provided for visualizing the whole proof. In other words, we provided to MDE practitioners the "glue" they need for connecting their domain with the logical domains needed for verification.

5.1 Heterogeneous Verification

Our problem is formally stated as a heterogeneous specification using CASL structuring constructs, with at least three logics: `CASL`, `CSMOF` and `QVTR`. We also perform logic translations through the implemented comorphisms which are `CSMOF2CASL` and `QVTR2CASL`. Next, there is an excerpt of the heterogeneous specification of the example.

```
(1)   logic CSMOF
      from QVTR/UML get UML |-> UMLMetamodel
      from QVTR/UML_WMult get UML |-> UMLConstraints

(2)   spec UMLProof = UMLMetamodel
      then %implies UMLConstraints end

(3)   logic QVTR
      from QVTR/uml2rdbms get uml2rdbms |-> QVTTransformation

(4)   logic CASL
      spec ModelTransformation = QVTTransformation with logic QVTR2CASL
      then %implies
        . key_RDBMS_Table
        . Top_PackageToSchema
        . Top_ClassToTable
      end
```

Within the `CSMOF` logic (1) we create two specifications from standard XMI files with the information of the class metamodel and SW-model in Figure 1. This implies the creation of a representation of signatures and formulas according to the institution defined in Section 3. Another specification is created (2) by extending `UMLMetamodel` and stating that `UMLConstraints` is implied. This means that every formula (multiplicity constraint) in the second specification can be derived, thus there must be a proof of it. This is how the satisfaction relation of the CSMOF institution is checked. We also use the `QVTR` logic (3) to create a specification from a standard `.qvt` file according to the institution defined in Section 3. The only difference with respect to the QVT standard is that instead of using OCL as the expressions language, we use for now a very simple language containing boolean connectives, the constants true and false, term equality, strings and variables. Finally, we move into CASL (through the comorphism `QVTR2CASL`) for creating another specification (4) in which the translation of key and rule formulas defined in Section 3 are implied by the transformation specification. When a proposition, e.g. `Top_ClassToTable`, is called from the CASL specification, a proof of the implication must be given. We can also translate our specifications and complement them with other constraints which cannot be stated as formulas of the former institutions. As an example we can state that there cannot be two `Classifier`s with the same name in the `UMLMetamodel` specification. For this purpose we are using the `CSMOF2CASL` comorphism as follows.

```
spec MoreProofs = UMLMetamodel with logic CSMOF2CASL
then %implies
  forall x,y : Classifier; str : String
  . name(x,str) /\ name(y,str) => x = y
end
```

Once our heterogeneous specification is processed, HETS constructs a development graph in which nodes correspond to specifications, some of them with open proof obligations, and arrows to dependencies between them. We have three proof obligations corresponding to those formulas marked as `%implies` within the specifications. Proof goals can be discharged using a logic-specific calculus, e.g. some prover for CASL in the example. The double arrows are heterogeneous theorem links, meaning that the logic changes along the arrow. In the example this corresponds to the extension of specifications by using the comorphisms. It can be noticed that we can use any other logic within the logics graph of HETS through comorphisms. This improves the proof capabilities of our environment.

5.2 Verification Properties

There are several properties that can be verified, some of them related to the computational nature of transformations and target properties of transformation languages, and other to the modeling nature of transformations [2]. The minimal requirement is conformance, i.e. that the source and target models (resp. the transformation specification) are syntactically well-formed instances of the source and target metamodels (resp. the transformation language). Our framework provides this verification in three parts. During the construction of CSMOF and QVTR theories, parsing and static analysis check whether signatures and formulas are well-formed, and (as we explained before) a SW-model within a signature is a structurally well-formed instance of the metamodel in the same signature, as well as a transformation specification given in a formula is well-formed with respect to the signature containing both source and target metamodels. Multiplicity constraints are verified when proving the satisfaction of CSMOF formulas. Finally, non-structural constraints are verified by extending both CSMOF and QVTR specifications using other logics, as CASL in the example. HETS also allows for disproving things using consistency checkers. This provides an additional point of view. In particular, we can check if a set of rules have contradictory conditions which could inhibit its execution.

In most cases a general-purpose logic, as provided by CASL, is enough to cover most of the verification approaches in [2]. The future inclusion of OCL as an institution will provide additional support in this sense. However, the verification process may depend on the problem to verify, since it is well-known that there is a "state explosion" problem when using automated checkers. Thus, automatic proofs are not always possible. In HETS it is possible to choose the tool we want to use. In this sense, we can choose not to use an automated theorem proving system, but for example an interactive theorem prover.

Verification interests go beyond these kinds of problems. When verifying a model transformation we want to consider its elements as a whole and not individually. In this sense, sometimes the notion of a transformation model is used, i.e. a model composed by the source and target metamodel, the transformation specification and the well-formedness rules. We have a transformation model in a QVTR theory (`QVTTransformation` in the example) which allows to add other properties by combining elements from the source and target metamodels and SW-models. With this we can state model syntax relations, trying to ensure that certain elements or structures of any input model will be transformed into other elements or structures in the output model. This problem arises when, for example, these relations cannot be inferred by just looking at the individual transformation rules. We can also state model semantics relations, e.g. temporal properties and refinement. Besides further work is needed to evaluate the alternatives, there are languages and tools, as ModalCASL and VSE (based on dynamic logics) commonly used for verifying these kinds of things. We could also be interested in working at another abstraction level, i.e. not considering specific SW-models but only metamodels and the transformation specification. This can be useful, for example, for proving that a transformation guarantees some model syntax relations when transforming any valid source SW-model. The problem here is that we need another institutional representation, e.g. we need to consider an abstract representation of a SW-model instead of a fixed one.

6 Related Work

There are some works that define environments for the comprehensive verification of MDE elements based on a unified mathematical formalism. As an example, in [13] rewriting logic is used to analyze MOF-like and QVT-like elements. Since rewriting logic was integrated into HETS [14], we can use these representations instead of using our comorphism into CASL. Nevertheless, since our institution is logic-independent it provides more flexibility for the definition of further specific comorphisms into other logics and languages (e.g. UML). In general, the use of a fixed unified mathematical formalism serving as a unique semantic basis can be quite restrictive. With our approach we can move between formalisms, and use a unified mathematical formalism if necessary (e.g. when transforming the whole specification into CASL).

In [15] the authors define a language-independent representation of metamodels and model transformations supporting many transformation languages. They also define mappings to the B and Z3 formalisms. Since they use only one generic language, only one semantic mapping needs to be defined for each target formalism. However, the semantic mapping should be semantics-preserving, and this aspect is not formally addressed in such work. In our case, comorphisms already preserve the semantics with respect to the satisfaction relation. Moreover, our comorphism into CASL and the corresponding implementation in HETS, provides the possibility of connecting our institutions to several logics and tools.

There are works representing the semantics of UML class diagrams with first-order logic, as in [16]. Since there are no so many alternatives for this representation, these works have similarities with our representation of extended CSMOF into CASL. In particular, the work in [16] is the nearest to ours from which we take many aspects, e.g. the "distinguishability" and "completeness of elements" axioms. In [17] the authors explain how class diagrams with OCL constraints can be translated into CASL. However, their definition is informally presented, and not in terms of a comorphism. In [18] the authors define a comorphism from UML class diagrams with rigidity constraints to ModalCASL (an extension of CASL). Since our $\mathcal{I}^{\mathrm{M}^+}$ institution is an adaptation of the institution for UML class diagrams, the comorphisms have some aspects in common, as the translation of formulas, but without the modal logic particularities.

Several approaches to heterogeneous specification have been developed for traditional software development, but there is little tool support. CafeOBJ [19] is a prominent approach based on the theory of institutions. However it provides a fixed cube of eight logics and twelve projections (formalized as institution morphisms), not allowing logic encodings (formalized as comorphisms). Thus, it is not an option for the definition of our environment. Moreover, HeteroGenius [20] is a framework, based on institutions, allowing the interaction between different external tools giving the user the possibility of performing hybrid analysis of a specification. However, the framework is not formally defined or available to be used as a basis for our environment.

7 Conclusions & Future Work

We have presented the implementation of an environment for the formal verification of different MDE aspects using heterogeneous verification approaches, which is based on a theoretical definition presented in a previous work [4]. The environment was integrated into HETS by defining comorphisms from institutions representing MDE elements to CASL, the core language of HETS. The existent connections between CASL and other logics within HETS broadens the spectrum of logical domains in which the verification of MDE elements can be addressed.

The environment supports a separation of duties between software developers (MDE and formal methods experts) such that a formal perspective is available whenever it is required. A developer can import the MDE elements, supplement this information with verification properties specified in other languages within the graph of logics supported by HETS, and perform the heterogeneous verification assisted by the tool. Since the implementation can generate a heterogeneous specification from the same files used by MDE practitioners, and there is no need of rewriting MDE building block in each logic involved, the environment is scalable without human assistance. Although our proposal is aligned with OMG standards, this idea can be potentially formalized for any transformation approach and language, which allows extending the approach as far as necessary. Finally, the environment is reliable since it is supported by a well-founded theory and by a mature tool in which there are several logics already defined.

Nevertheless, we still have some open issues. A current drawback is the inexistence of an institution for OCL which is a language in which QVT is strongly based. For now we have considered a very simple expressions language, but the definition of an institution for OCL is subject of further work. In the same sense, we expect to extend the institutions to include some elements not considered and give them tool support, besides exploring other options for the verification of transformation properties. This will strengthen the formal environment for MDE. Since our institutions formalize languages strongly related with those in the UML ecosystem, it will be interesting to explore the possibility of integrating them with other languages, as those already defined as institutions in [21].

We need to continue bridging the gap between MDE and formal verification in terms of tool development in order to practitioners really be able to benefit from our approach. First, we can connect the definition of the MDE elements in any popular tool with an automatic generation of the heterogeneous specification, as explained in Section 5.1, and the execution of Hets using this specification. Moreover, we could perform an automated verification of some properties (if possible) by running Hets in the background and providing a better user interface to show the problems found by Hets. For this to be possible, we need to improve feedback from existing formal tools. This needs better traceability between the problem definition and the results given by a verification tool. We can define some traceability links from comorphisms, interpret the output of the verification tool and return something that the MDE practitioner can interpret. This interpretation is like defining a transformation between the domain of outputs of the verification tool and the domain of messages in MDE.

Moreover, as described in Section 5.2, the environment deals with many verification properties, but a deeper understanding of this (as for example about the behavior of models) is a must. In this sense, we can use the knowledge in [2] to provide a guide for the selection of the "right" verification approach for the problem which is of interest to verify. We also need to apply our approach to industrial, real-size examples for strengthening the results.

A final topic of interest, somewhat related to this work, is to explore the possibility of leveraging the capabilities of Hets by using MDE elements as a metalanguage for expressing logics and comorphisms. If metamodels are defined for different logics, model transformations can be used to express comorphisms between them. Models can represent specifications within corresponding logics and an automatic process can generate their representation to the Hets engine. This could eventually simplify the definition of logics and comorphisms.

References

1. Kent, S.: Model driven engineering. LNCS, vol. 2335, pp. 286–298. Springer (2002)
2. Calegari, D., Szasz, N.: Verification of model transformations: A survey of the state-of-the-art. ENTCS, vol. 292, pp. 5–25. Elsevier (2013)
3. Mossakowski, T.: Heterogeneous specification and the Heterogeneous Tool Set. Technical report, Universitaet Bremen (2005) Habilitation thesis.

4. Calegari, D., Szasz, N.: Institution-based semantics for MOF and QVT-relations. LNCS, vol. 8195, pp. 34–50. Springer (2013)
5. Goguen, J.A., Burstall, R.M.: Institutions: Abstract model theory for specification and programming. Journal of the ACM **39** (1992) 95–146
6. Mossakowski, T., Maeder, C., Lüttich, K.: The Heterogeneous Tool Set. LNCS, vol. 4424, pp. 519–522. Springer (2007)
7. Mossakowski, T., Haxthausen, A.E., Sannella, D., Tarlecki, A.: CASL- the common algebraic specification language: Semantics and proof theory. Computers and Artificial Intelligence **22** (2003) 285–321
8. OMG: Meta Object Facility (MOF) 2.0 Core Specification. Specification Version 2.0, Object Management Group (2003)
9. OMG: Object Constraint Language. Formal Specification Version 2.2, Object Management Group (2010)
10. OMG: Meta Object Facility (MOF) 2.0 Query/View/Transformation. Final Adopted Specification Version 1.1, Object Management Group (2009)
11. Calegari, D.: Heterogeneous Verification of Model Transformations. PhD thesis, Universidad de la República - PEDECIBA (2014) url: https://www.fing.edu.uy/inco/pedeciba/bibliote/tesis/tesisd-calegari.pdf
12. Codescu, M.: Generalized theoroidal institution comorphisms. LNCS, vol. 5486, pp. 88–101. Springer (2008)
13. Boronat, A., Heckel, R., Meseguer, J.: Rewriting logic semantics and verification of model transformations. LNCS, vol. 5503, pp. 18–33. Springer (2009)
14. Codescu, M., Mossakowski, T., Riesco, A., Maeder, C.: Integrating Maude into HETS. LNCS, vol. 6486, pp. 60–75. Springer (2011)
15. Lano, K., Rahimi, S.K.: Model transformation specification and design. Advances in Computers **85** (2012) 123–163
16. Shan, L., Zhu, H.: Semantics of metamodels in UML. In: Proc. TASE, IEEE Computer Society (2009) 55–62
17. Bidoit, M., Hennicker, R., Tort, F., Wirsing, M.: Correct realizations of interface constraints with OCL. LNCS, vol. 1723, pp. 399–415. Springer (1999)
18. James, P., Knapp, A., Mossakowski, T., Roggenbach, M.: Designing domain specific languages: A craftsman's approach for the railway domain using CASL. LNCS, vol. 7841, pp. 178–194. Springer (2013)
19. Diaconescu, R., Futatsugi, K.: CafeOBJ Report: The Language, Proof Techniques, and Methodologies for Object-Oriented Algebraic Specification. Volume 6 of AMAST Series in Computing. World Scientific (1998)
20. Giménez, M., Moscato, M., López, C., Frias, M.: Heterogenius: A framework for hybrid analysis of heterogeneous software specifications. EPTCS, vol. 139, pp. 65–70. (2018)
21. Cengarle, M.V., Knapp, A., Tarlecki, A., Wirsing, M.: A Heterogeneous Approach to UML Semantics LNCS, vol. 5065, pp. 383–402. Springer (2008)

A Proposal for Integrating Formal Methods into a Lightweight UML-driven Development Process

Thiago C. de Sousa[1] and Paulo Sérgio Muniz Silva[2]

[1] State University of Piauí, Brazil
thiago@uespi.br
[2] University of São Paulo, Brazil
paulo.muniz@usp.br

Abstract. The best practices of software engineering indicate that the verification activity is essential to achieve some quality during the software construction. In UML-based development processes, one of their main focuses is the detection of inconsistencies in diagrams that represent the software. However, most of these processes, such as ICONIX, apply only informal techniques (eg. visual model inspection), often implying the negligence of that activity by developers. Moreover, with the advance of automated verification tools, formal methods, such as Event-B, are increasingly attracting the attention of software companies. However, it is still difficult to convince developers to adopt them, because they are not acquainted with some of their mathematical concepts. Thus, this paper presents a proposal for the inclusion of Event-B within ICONIX, giving rise to BICONIX, an object-oriented development process that supports inconsistencies formal verification. Specifically, this work shows how this merger can assist the verification activity in well-defined check points of the proposed process.

Keywords: Formal Verification, UML, ICONIX, Event-B

1 Introduction

UML has become the "de facto" standard for software modeling and, nowadays, there are a lot of development processes that use UML diagrams to create partial models of the system being produced. These models usually describe a system from different viewpoints and levels of abstraction and, frequently, lead to a number of inconsistency issues, which are well-known by the software engineering community. One of these problems is to make sure that all models respect the constraints (business rules and functional properties) imposed by the application domain and/or by the stakeholders. Another issue is how to ensure that each software model has a unique interpretation (precise semantics), which means that it cannot be understood in different ways by two or more developers. Finally, there is the problem of checking whether the semantics of an abstract model is preserved by its detailed versions after one or more successive refinements.

On the one hand, most of software development processes based on UML, such as RUP and ICONIX [1], include the verification of these inconsistencies as

an essential task. However, many verification techniques are based on inspections of the models and, as the UML models and the constraints are expressed by informal languages, the inspections are usually carried out manually and visually, making them costly and strongly dependent on the skill and experience of the developer.

On the other hand, formal methods, such as VDM [2], B [3], Event-B [4] and Z [5], have supported not only to check model inconsistencies precisely and automatically, but also work as software specification and modeling tools. However, despite the effectiveness of formal techniques in filling the gap between requirements specification and implementation and ensuring system correctness by construction (verification mechanism), industrial practitioners are still reluctant to fully adopt them.

As we can realize, although inconsistencies verification is a critical issue in software engineering, the widely used software development processes based on UML have no efficient mechanism to perform this task. It is also observed that formal methods provide effective techniques to address this problem, but they do not attract the attention of the community. So, for many industrial practitioners it would be very useful to have a well-known UML-based process with support for formal verification. In this work we present an approach for the integration of the ICONIX process with the Event-B formal method. More precisely, we show how the Event-B formalism can be incorporated into the ICONIX steps/stages in order to provide a mechanism to check for inconsistencies.

In the next section we present some theoretical fundamentals, introducing the features of the ICONIX process, as well as the main concepts of the Event-B language. In section 3 we explain our approach named BICONIX in more details, showing an structure overview and the main tasks of each phase. In section 4, we show some related work, and finally we reserve the last section for further discussions about how BICONIX can assist the verification task and directions for future work.

2 Background

The ICONIX process can be considered a pure, lightweight and practical, but also powerful methodology. ICONIX is not as bureaucratic as RUP, which means that it does not generate a lot of documentation. And despite being a simple process like XP (eXtreme Programming), it does not lack the Analysis and Design phase, providing a very simple step-by-step guiding rules during the whole process. Moreover, ICONIX uses only four diagrams (use cases, robustness, sequence and class), follows the iterative and incremental development cycle and brings the developer to a “mandatory” and well-defined verification task among its phases in order to check requirements compliance.

Recently, a variant of the B formal method has been successfully applied in some companies from different fields, such as aerospace, services and railways. Event-B is a state-based formal method for modeling systems based on predicate logic and set theory where the refinement mechanism, which allow us to build a

model gradually by making it more and more precise (that is, from an abstract model to a concrete one), and the consistency checking, which allows us to verify the validity of the properties of a system, are guaranteed by mathematical proof obligations. These features have been supported by an open-source platform (based on Eclipse IDE) named Rodin, that is constantly improved and extended by the community via plugins.

3 The BICONIX Process

In this section we present an object-oriented development process named BICONIX that support supports formal verification of inconsistencies. The proposed process keeps basically the main characteristics of ICONIX and extends it by incorporating a specific role (the Event-B Profile) for allowing the addition of invariants (business rules and functional properties), guards, actions and refinement relations among the models generated automatically in each of its first three phases. An overview of BICONIX can be seen in Figure 1.

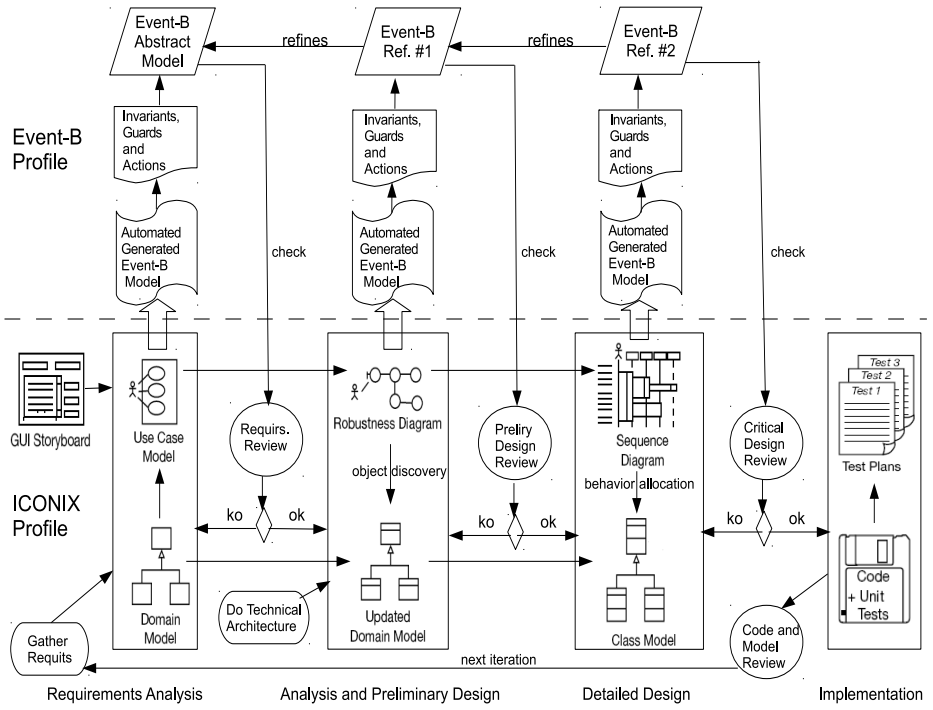

Fig. 1. Overview of the BICONIX process

The BICONIX process has two dimensions: the horizontal axis represents the temporal order and shows the lifecycle aspects of the development process; the vertical axis is used to represent both structural and behavioral software aspects. The first dimension represents the dynamic aspect of the process and is expressed in terms of phases, milestones and iterations. The second dimension is the static aspect of the process, as it is described in terms of models, activities, workflows, artifacts and profiles.

The BICONIX process is purposely very similar to ICONIX in order to invite regular developers to using it. So, BICONIX has also four sequential phases, each one concluded by a milestone. At the end of a phase there is an execution of a critical review in order to determine whether the objectives have been achieved. A positive evaluation allows the project to move forward to the next stage. Basically, the artifacts, activities and workflows of the BICONIX process are the same ones of ICONIX, only including those ones related to the Event-B profile, keeping the lightweight aspect of ICONIX. Due to space limitations, we will not detail BICONIX, only focusing on the differences from ICONIX.

One of the main novelties of BICONIX compared to the original process is the support of formal verification, which appears after the transition between the ICONIX and Event-B profiles. At the end of each of its first three phases, the Event-B specialist receives the diagrams produced by the ICONIX developer and uses the Rodin platform, extended with some transformation rules, for automatic translation of them into the Event-B language. Basically, the artifacts that represent the static part of the software (Domain Model, Updated Domain Model and Class Diagram) generate sets, relations and type invariants in Event-B, and the artifacts that represent the dynamic part of the software (Use Case, Robustness and Sequence Diagrams) are mapped to events.

The generated formal model may be augmented with constraints (invariants) and pre/post-conditions (guards/actions). In the first phase (Requirements Analysis), the Event-B expert can define constraints over the domain model and pre/post-conditions on the use cases. In the second phase (Analysis and Preliminary Design), some constraints over the updated domain model (with attributes) and pre/post-conditions on the robustness diagrams can be described. Finally, at the stage of Detailed Design, the specialist can include constraints over the class model (with methods) and pre/post-conditions on the sequence diagrams.

The final formal model produced is then checked automatically by the tool in order to detect errors. If there are any issues, they are discussed by the two roles (profiles) during the review activity of each phase (requirements review, preliminary design review or critical design review). Essentially, the problems occur due to invariants broken by event actions or due to a detailed model that is not following the formal refinement rules. For mapping back the Event-B model errors as ICONIX artifacts issues, the specialists are aided by an implicit dictionary (e.g. the constraint word in the ICONIX context is synonyms of the invariant word in the context of Event-B). So, the developers must decide which steps of the current phase, including the diagram construction, should be redone in order to correct the errors, before proceeding to the next one.

4 Related Work

Relevant work present formal UML-based methods. Runde *et al.* [6] present a formal method called STAIRS based on Interactions and Sequence diagrams as defined in UML 2.0 , with the horizontal and vertical refinement notions. Ke *et al.* [7] detail a sound process for developing critical systems named rCOS, which translates some UML diagrams into a formal language and checks them using the FDR tool. Ahrendt *et al.* [8] present a formal object-oriented process named KeY, which uses UML diagrams and OCL annotations to generate verified code in Java Card format using a tool developed by the authors. However, they are not based on any known methodology, which may not encourage their use.

Other work use our same approach to provide a mapping from UML to a well known formal language. Laleau and Mammar [9] provide a method to generate B specifications from UML (class, state and collaboration diagrams) and check them using the AtelierB tool. Lausdahl *et al.* [10] present an approach to translate Class and Sequence diagrams to VDM++ and check them using the Overture tool. Miao *et al.* [11] show how to translate Class, Sequence and Statecharts diagrams into Object-Z and check them using the OZRose tool. But there is no mention to integrate them in a diffused methodology like we do.

5 Discussions and Future Work

In this paper we have proposed an approach for including a formal method (Event-B) into a lightweight UML-based methodology (ICONIX) in order to aid industrial practitioners with the verification task. Since there is the formalization of the artifacts, the user of the BICONIX process has confidence to check both syntactic (eg. every use case must have a name) and semantic (eg. no cycles in a use case diagram) modeling issues accurately.

We believe that the introduction of a formal method in a development process brings, under the technical point of view, two major benefits. The first one is the discovery of modeling problems at the early stages, which contributes significantly to the reduction of rework, essential to minimize the technical debt. A second gain would be the obligation of performing the verification activity in specific check points of the process, since BICONIX forces the execution of this task between the transition of its phases, thus contributing to the dissemination of the formal analysis culture.

From the manager's viewpoint, it is important to emphasize that the BICONIX process proposes the coexistence of two technical roles to perform its many activities: the ICONIX and the Event-B profiles. At a glance, this feature may be considered an additional issue to manager control, impacting negatively upon the management of projects that follow the process. However, we believe that the inclusion of a specialist that fits the Event-B profile, for managers who already have experience in the ICONIX process, should not be a major issue, since the tasks under the liability of the expert are performed only at the end of each phase, having no impact on the ICONIX process kernel.

Finally, this work requires important future improvements, many of them related to the BICONIX limitations: the inclusion of elements in the Event-B language for supporting the object-oriented semantics; the addition of other standard elements of the UML diagrams (eg. messages types of Sequence Diagram); the definition of refinement patterns to accelerate the description of *gluing invariants* among the generated Event-B models in each phase; the complementation of the process to enable the addition of invariants and guards/actions in the Implementation phase, providing an integration with a language that allows specifying these details directly in the source code; the integration of an easier constraint language in order to reduce the responsibilities of Event-B specialist; the association of a tool for automatic test generation in order to reinforce the verification mechanism; the incorporation of a mechanism to assist the transition from requirements specification to Use Cases and Domain models, developing, for example, a controlled language to express them; the improvement of the feedback from Event-B errors provided by the Rodin platform to diagrams issues; and the elaboration of a controlled experiment to assess the viability of the proposed process in real projects.

References

1. Rosenberg, D., Stephens, M.: Use Case Driven Object Modeling with UML: Theory and Practice. Apress (2007)
2. Bjørner, D., Jones, C.B.: The Vienna Development Method: The Meta-Language. Lecture Notes in Computer Science **61** (1978)
3. Abrial, J.R.: The B-book: assigning programs to meanings. 1st edn. Cambridge University Press, New York, NY, USA (1996)
4. Abrial, J.R.: Modeling in Event-B: System and Software Engineering. 1st edn. Cambridge University Press, New York, NY, USA (2010)
5. Spivey, J.: The Z notation - a reference manual. Prentice Hall International Series in Computer Science (1989) I–XI, 1–155
6. Runde, R.K., Haugen, O., Stolen, K.: The Pragmatics of STAIRS. In de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.P., eds.: Proceedings of the International Symposia on Formal Methods for Components and Objects (FMCO). Volume 4111 of Lecture Notes in Computer Science., Springer (2005) 88–114
7. Ke, W., Li, X., Liu, Z., Stolz, V.: rcos: a formal model-driven engineering method for component-based software. Frontiers of Computer Science **6**(1) (2012) 17–39
8. Ahrendt, W., Beckert, B., Hähnle, R., Schmitt, P.H.: KeY: A formal method for object-oriented systems. In: Procs. of 9th. Intl. Conf. on Formal Methods for Open Object-Based Distributed Systems, Cyprus, 2007. LNCS, Springer (2007)
9. Laleau, R., Mammar, A.: An Overview of a Method and its Support Tool for Generating B Specifications from UML Notations. In: Proceedings of the IEEE International Conference on Automated Software Engineering (ASE '00), Washington, DC, USA, IEEE Computer Society (2000) 269–
10. Lausdahl, K., Lintrup, H., Larsen, P.: Connecting UML and VDM++ with open tool support. In: FM 2009: Formal Methods. Volume 5850 of LNCS. Springer (2009)
11. Miao, H., Liu, L., Li, L.: Formalizing UML Models with Object-Z. In: Formal Methods and Software Engineering. Volume 2495 of LNCS. Springer (2002)

Including Running System Implementations in the Simulation of System of Systems Models

Kenneth Lausdahl[1],
Claus Ballegård Nielsen[1], and Klaus Kristensen[2]

[1] Department of Engineering, Aarhus University.
Finlandsgade 22, Aarhus N 8200 Denmark
{lausdahl,clausbn}@eng.au.dk
[2] Bang & Olufsen, Denmark
krt@bang-olufsen.dk

Abstract. The formal modelling of System of Systems is challenged by the autonomy of the participating constituent systems as it may not be possible to obtain the implementation details needed to create a descriptive model. This paper describes an extension for the Symphony tool that enables formal models of System of Systems to be connected and simulated with externally running system implementations. An industrial case study of a Bang & Olufsen sound system is used to show the application and feasibility of the approach.

Keywords: System of Systems, Formal Modelling, External Systems

1 Introduction

A System of Systems (SoS) is a type of system that is composed of largely independent constituent systems, which collaborate in order to reach a common goal [4]. In their nature SoSs consist of distributed constituent systems that have a high degree of autonomy and often are developed and evolved individually. The complex structures and interactions found in an SoS can be described formally, and the challenges faced by system developers can be analysed through formal modelling [3].

The COMPASS Modelling Language (CML) is a formal modelling notation aimed at creating models of SoS architectures and SoS properties in order to help developers analysing SoS [8]. Despite the tool-support and a formalism aimed at describing SoS, the modelling of SoS is however still challenging. The two main causes of this is that: 1) the independent owners of the systems may not be willing to share all knowledge on their system implementation, making it difficult to create models with the correct behaviour, 2) the systems may be legacy systems with no precise description of their internals. In this paper, we show how SoS models can incorporate external constituent systems as part of the model simulation through improved tool-support.

The remainder of this paper is structured as follows: Section 2 gives an overview of the challenges in formal modelling of SoS. The principles of the simulation between the CML simulator and the external system, as well as details of the tool extension are described in Section 3. A small study on using the external simulation in an industrial case is presented in Section 4, before we draw conclusions in Section 5.

2 System of Systems Engineering and Formal Modelling

SoS Engineering is challenged by a strong degree of autonomy, evolution and emergence [4]. In order to address this, the use of formal methods has been proposed as a way of providing better analyses of the system design [1]. The formal languages directly aimed at SoS modelling are however still challenged by the autonomy constituent systems. As there can be conflicting interests between the owners of the individual constituent systems, or because parts of the SoS are delivered by a supplier that has no interest in the SoS, it may be difficult to create sufficiently descriptive formal models. The owners of the independent constituent systems may not be willing to share implementation details for competitive reasons, and suppliers may just have delivered a commercial off-the-shelf product for which they have no interest in delivering documentation on its internals. Finally, some of the constituent systems in the SoS maybe legacy systems for which documentation of their system implementation is no longer producible. This creates challenges in the simulation of the executable formal models, as a meaningful simulation depends on the actual behaviour of the systems involved.

We present a tool extension to simulate SoS models for which the behaviour of some of the constituent systems can only be obtained via the actual running system, by establishing connectivity between the simulator and an external system.

Including running systems in the simulation is already possible in other formal methods. The Overture tool has a mechanism that allows VDM models to delegate calls to external Java libraries [5], and for Coloured Petri Nets, CPN Tools offers the Access/CPN functionality which enables the integration of CPN models with external applications via a TCP/IP stream [7]. These approaches do, however, use a different approach that either requires modification of the model or for the external system to control the simulator.

2.1 The COMPASS Modelling Language

CML combines elements from the VDM state-based formal method [2] and the CSP process-based formal method [6] in order to express both the structure and behaviour of an SoS. An SoS is modelled as a collection of constituent systems with their behaviour specified in processes, which have communication channels defined between them.

The CML language is built around collections of types, values, functions, operations, classes, processes and channels, with the types, values, functions and classes originating from VDM, while channels are taken from CSP/Circus. The processes are the central constructs in CML as they act as the constituents in the SoS. A process can define and maintain state using the VDM based type system. Finally, actions are used to express the reactive behaviour of a process. Channels are globally defined and are used for defining the communication and synchronization events between processes.

The Symphony tool[3] is an Eclipse-based IDE that provides a CML parser, typechecker and interpreter & debugger for model simulation, as well as range of static analyses such as a theorem prover, proof obligation generator and model checking.

An example of a very basic CML model is given in Listing 1.1, which shows the definition of a channel c carrying a nat type, and three small systems represented by the processes P, A and B. A parallel composition is made between the A and B processes

[3] Symphony Tool webpage: https://github.com/symphonytool/

as defined by the process P. A process, such as P, that defines the system composition at the highest level is known as the *top* process. When making a composition, a range of channels can be given to denote the channels that the composed process can synchronize on, based on the events occurring on the channels. In the given example the composition defines that the processes will synchronise on events occurring on channel c. Process A and B have a very simple functionality, where A is a system that defines an action (indicated by @) to synchronize the value 1 on channel c, while B is a system that synchronizes a value placed on channel c (indicated by ?) and assigns it to x.

Within the simulator, a CML model is represented as a tree structure of model behaviours. The *top* process is used to define the entry point for the simulation, and the simulation is performed by moving from the top and down the tree to inspect and execute underlying behaviours. An inspection will reveal the collection of events that the composed systems can execute. The structure for Listing 1.1 is shown in Figure 1, where the inspection of the composition results in the collection of possible events. A parent in the tree structure functions as a *coordinator* for its children, meaning that it will filter the collection of events available to the children and thereby control the possible execution.

```
channels c : nat
process P = A [|{c}|] B

process A = begin
  @ c.1 -> Skip
end

process B = begin
  @ c?x -> Skip
end
```

Listing 1.1: CML example

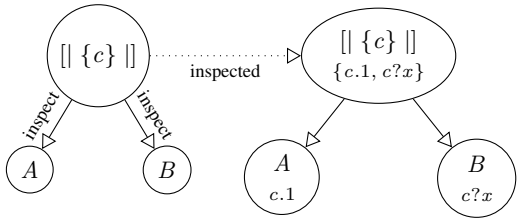

Fig. 1: Internal simulator representation of P from Listing 1.1.

3 Tool Support for Simulation with External Systems

The CML simulator in the Symphony tool has been extended such that systems that are running externally to the simulation environment can be included in the simulation of CML models. This means that concrete system implementations can be run and executed in parallel to the simulation of the model, allowing the behaviour of both the model and the system to affect each other. The presented approach can be divided into three subsections: the principles of the approach (Section 3.1), the changes to the simulator (Section 3.3) and the protocol used to establish the interaction (Section 3.2).

3.1 Principles

When an external running system implementation is going to be part of the simulation of a CML model, it will act as one of the constituent systems defined in the overall SoS. As each constituent system is expected to have its own specified behaviour and have a specified way of communicating with the other constituent systems in the SoS model, it is possible to make the external system act as a constituent system in the model. Essentially, the tool extension will replace a process definition in the model with a skeleton that will delegate the interaction to the external system.

The Symphony tool will be in control of the simulation and will act as a *coordinator* in the simulation, while the external system will act as a client. The coordinator is responsible for starting the simulation and for controlling the flow of the simulation, once started. A special debug configuration in the Symphony tool is used to setup the simulator as a coordinator and for interacting with external systems. In the configuration of the coordinator the top level process is selected, as well as the processes that will be handled by an external system. The debug configuration also requires address and port information to be specified, such that the clients can connect to the coordinator.

During the simulation the coordinator will control the flow of the model execution according to the model specification. The possible transitions that can be made in the model are computed in the same way as for a normal simulator. The only change is that some of the computations are performed by external systems. From a simulation point of view the executing model is essentially seen as one model, which is not distinguishable from a normal simulation.

3.2 Protocol

The connectivity and data exchange between the simulator and the externally running system is performed via a custom protocol built on top of a TCP connection. The protocol for handling the simulation reflects the simulator's way of inspecting and execution behaviours, and the protocol has been formally specified in a CML model. To facilitate the interoperability between the simulator's Java implementation and the running systems, that may be software implemented with different technologies on different platforms, a data-interchange format is used. The JavaScript Object Notation (JSON) is used, as this will ensure interoperability. The protocol mirrors the execution flow of the normal simulator by providing message types for the inspection and execution of behaviours. The protocol keeps a flow state in order to ensure that an execution message is only allowed if an inspection message has previously been processed. The protocol also contains message types that allow clients to register on the coordinator, as well as messages for disconnecting. The network functions as client–server relationship, with the coordinator functioning as a server that the clients can connect to during the initialisation of the simulation.

3.3 Simulation with External System

The Symphony simulator has been extended such that it changes the way it handles behaviours. Essentially, all the processes that have been specified as being handled by an external system in the debug configuration will be replaced with a skeleton that delegates the processing to the external system via a network setup. As the protocol imitates the inspection and execution calls occurring in the normal simulation, the coordinator can use most of the existing simulator implementation and function almost like a normal simulator, just with some calls being delegated over the network.

This is illustrated in Figure 2, that shows: (a) the tree structure of a normal simulator and (b) the same tree structure for simulation where the process B is an external system. The dashed line represents the network connection between the simulators.

The external system needs to have an adaptor between the system implementation and the network setup to the simulator. The adaptor is responsible for connecting to the

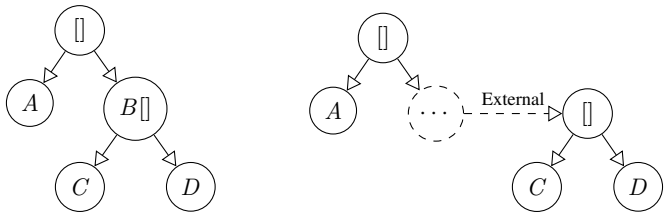

(a) Standard Simulator (b) Simulation with External System

Fig. 2: Illustration of the internal simulator change.

simulator over the network and for implementing the simulation protocol specifying the possible interactions (Section 3.2). The implementation of the protocol involves two steps: 1) a mapping of the types of the concrete programming language to the CML types described in the simulation protocol, and 2) a mapping from the protocol messages to operations in the external system itself. Essentially, the external system needs to provide the simulator with a list of transitions that represents the possible events that the system can perform given the current system state, and allow for the simulator to send an execute message that can change the system state. This can be implemented through the use of a state machine in the adaptor, which is used in the case study Section 4.

4 Case Study

The approach was examined through an industrial case study involving a Bang & Olufsen (B&O) home Audio/Video (A/V) network for connecting devices (such as audio, video and legacy audio products) distributed across a user's home. These devices may be produced by competing manufacturers, but need to interact in order to deliver a service to the user. As such, the A/V network forms an SoS of devices that can both be heterogeneous and legacy systems. The devices may provide stand-alone streaming or media content rendering services, but the SoS needs to deliver a coherent experience. A key part of the system is the A/V control that has the task of managing which of the devices should stream and render media. This behaviour is defined through streaming contracts that specify the interactions between the devices.

A CML model of the A/V network streaming contracts is used to examine the presented approach. The CML model contains 2 devices: (1) the AV control representing the user's interaction semantics, and (2) a streaming device which renders media on the basis of control events from the A/V control. The CML model describes the streaming contracts by defining a set of semantically defined transition rules for distributed state synchronization and distributed operation calls.

The externally running system used for this case study is a B&O developed C++ implementation of the streaming contract. The implementation makes use of two interfaces: 1) a call-back interface where application layer clients will be notified of transition operations and states changes, and 2) an interface that functions as an application layer abstraction creating an adaptor towards the platform specific streaming implementations. These two interfaces maps to the channels defined in the CML model of the streaming contract, and the semantics of these interfaces' implementations satisfy the transition rules defined for the streaming contract. In order to map the externally running system to the events being described in CML, a state machine has been implemented where the channel events of the CML model are translated into operations

calls in the C++ implementation. For the network setup between the simulator and the external system some glue code in the C++ implementation is used to handle the mapping for states, types and operation logic. This also works as a wrapper that creates the needed JSON messages for the external system.

Being able to include the running system as part of the model simulation, enabled B&O to check the behaviour described in a model against the real system implementation. In the SoS model a process modelling a constituent system was replaced with the actual system being modelled. Testing a concrete system up against a model showed that the behaviour described in the model corresponded to the systems implementation for the given modelled scenario, and vice versa. This same approach can be used to replace parts of the SoS model with existing systems from other manufacturers or with legacy systems.

5 Conclusions

We have presented an approach for including externally running systems in the simulation of formal models of SoSs. By the use of a tool extension it has become possible to include constituent systems that, for various reasons, cannot be modelled into the simulation of a complete SoS model. We have shown the application of the approach through a small industrial case study and have demonstrated its feasibility. Allowing the simulation of SoS partly consisting of running constituent systems enables developers to address the challenges of autonomy and legacy often found in SoS development, while still being able to use a formalism targeted SoS development. We believe that the approach described in this paper can be used as an inspiration for tool builders of other formal methods that have tool–support for performing simulations of models.

Acknowledgment

The work presented here is partly supported by the EU Framework 7 Integrated Project: COMPASS, Grant Agreement 287829. The authors would also like to thank the anonymous reviewers for their valuable comments and suggestions.

References

1. D. Drusinksy and M.-T. Shing. Creation and Evaluation of Formal Specifications for System-of-Systems Development. In *IEEE SMC 2005*, pages 1864 – 1869 Vol. 2, oct. 2005.
2. John Fitzgerald and Peter Gorm Larsen. *Modelling Systems – Practical Tools and Techniques in Software Development*. Cambridge University Press, Second edition, 2009.
3. John Fitzgerald, Peter Gorm Larsen, and Jim Woodcock. Modelling and Analysis Technology for Systems of Systems Engineering: Research Challenges. In *INCOSE*, July 2012.
4. Mark W Maier. Integrated modeling: A unified approach to system engineering. *Journal of Systems and Software*, 32(2):101–119, 1996.
5. Claus Ballegaard Nielsen, Kenneth Lausdahl, and Peter Gorm Larsen. Combining VDM with Executable Code. In *ABZ2012*, volume 7316 of *LNCS*, pages 266–279, 2012.
6. A. W. Roscoe. *Understanding Concurrent Systems*. Springer, 2010.
7. Michael Westergaard and Lars Kristensen. The access/cpn framework: A tool for interacting with the cpn tools simulator. In *PETRI NETS 09*, pages 313–322, 2009.
8. J. Woodcock, A. Cavalcanti, J. Fitzgerald, P. Larsen, A. Miyazawa, and S. Perry. Features of CML: a Formal Modelling Language for Systems of Systems. In *SoSE 2012*. IEEE, July 2012.

Purification of ESTEREL Programs *

Nir Koblenc[1] and Shmuel Tyszberowicz[1,2]

[1] Department of Mathematics and Computer Science,
Open University of Israel, Raanana, Israel
[2] School of Computer Science,
Academic College of Tel-Aviv Yaffo, Tel-Aviv, Israel
skoblenc@gmail.com, tyshbe@tau.ac.il

Abstract. ESTEREL is a synchronous programming language dedicated to control-dominated reactive systems. XEVE, an ESTEREL verification tool, verifies circuit descriptions generated from the source programs. However, the ESTEREL compiler produces circuits with a behavior equivalent to the original program only for programs that do not handle data. The abstraction performed by the compiler removes the data, yielding an over-approximation that might violate safety properties even when the source program does not, causing XEVE to reject correct programs. We introduce an automatic abstraction process for ESTEREL programs developed to tackle this problem. When applied to a program, it results in a pure program with an equal external observable behavior. When applying the abstraction to a program augmented with observers that monitor safety properties, the ESTEREL compiler can compile the resulting pure program into a verifiable equivalent circuit. We have built a prototype tool that implements the abstraction and used it to purify control programs and robotic systems.

Keywords: Verification, Abstraction, Reactive Systems, Esterel

1 Introduction

Reactive systems are computer systems that continuously react to their environment at a speed determined by the environment. Most industrial real-time systems are reactive [6]. ESTEREL[1] is an imperative concurrent language for the development of industrial-strength reactive systems, which is especially well suited for control-dominated reactive systems such as real-time process control systems, embedded systems, communication protocols, peripheral drivers, human-machine interfaces, and others [1]. ESTEREL belongs to the family of

* The paper is based on the master's thesis of the first author. A draft of the thesis and the prototype tool are available at http://www.cs.tau.ac.il/~tyshbe/NIR/nirThesisDraft.html. This work has been partially supported by GIF (grant No. 1131-9.6/2011).

[1] We refer to ESTEREL v5.92, which we use for teaching. The toolset and the documentation are available at http://www-sop.inria.fr/esterel.org/filesv5_92/.

synchronous languages— languages that are based on the synchrony hypothesis, which states that a program instantaneously reacts to its input. Control is assumed to take no time and thus output is broadcast right when the input arrives. The notion of simultaneity is captured by the concept of *event*, which is a set of simultaneous occurrence of (possibly valued) signals [6]. A full definition of the language, as of version 5.91, can be found in [1].

Reactive systems are often used to control safety-critical systems. They therefore require rigorous design methods, and formal verification must be considered [6]. Constructing and verifying a formal specification of the system is very potent in detecting errors already at the formal specification development phase [9]. However, there still is a risk that there would be inconsistencies between the formal specification and the eventual product. Even while the specification is formally verified, the product itself may still be erroneous, and we want to verify that the program satisfies its safety requirements.

Verification by observers [3] is an approach to verify code. *Observers* are program modules monitoring the program, testing that a property is satisfied and broadcasting specific signals when the property is violated. The observers are composed in parallel to the original program, and the resulting program is compiled using an ESTEREL compiler into a finite automaton. The properties of the automaton are verified using tools such as the X ESTEREL VERIFICATION ENVIRONMENT (XEVE) [3], reducing the verification problem to reachability problem in finite automata – finding if there exists an execution trace from the initial state to a state emitting one or more of those special observer signals.

The XEVE verification environment requires the program to be compiled into Berkeley Logic Interchange Format (BLIF), a logic-level hardware hierarchical circuit description in a textual form; however, ESTEREL compiler, as of version 5.91[2], can either compile pure ESTEREL programs into BLIF files without changing their semantics, or, using the `-soft` option, abstract the data and compile only the control aspect into ESTEREL [2, Section 2.3.4]. Pure ESTEREL programs only handle pure signals, i.e., they involve no valued signals, types, constants, functions, procedures, tasks, or variables [2]. In this work we collectively refer to valued signals, sensors, and variables as *valued objects*. The problem with compiling a non-pure program into a BLIF circuit using the `-soft` option is that programs whose correctness depends on run-time values might be disqualified by XEVE as "unsafe".

Yet many control schemes, such as signal processors and closed-loop feedback controllers, receive numerical inputs, conduct numerical calculations, and emit numerical outputs. We *purify* such programs to allow automatic verification of their properties. This is a transformation of ESTEREL programs handling valued objects into Pure ESTEREL programs. It abstracts an unbounded, concrete system that handles data by replacing objects that take values from theoretically-infinite domains with pure signals to receive a finite system. The abstraction

[2] We chose to focus on ESTEREL v5 since it is free, suitable for teaching, and can be verified using the free XEVE tool which is part of the ESTEREL v5_92 distribution, whereas ESTEREL v7 [10] is commercial. XEVE is used in the industry [3].

preserves the external observable behavior of the original program, i.e. there is a correspondence in terms of inputs, outputs, and timings between the two programs. Usually, abstraction-based proof techniques are sound but not complete, since the abstraction is done such that every property proven to be satisfied by the abstract system has a concrete version which holds on the concrete system, yet the other way around is unnecessarily true [4]. However, verification using our abstraction technique is both sound and complete, since the abstract program adds no new behaviors to those displayed by the concrete program; in particular every pure signal is emitted by the observer-augmented concrete program if and only if it is emitted by its purified version. The abstraction alters the semantics of a few operations; however, these changes remain internal, i.e. the interaction with the environment is unaffected.

We suggest an algorithm that automatically purifies a certain group of programs and we characterize the class of programs verifiable using our technique. The main contribution of this work is extending the class of programs verifiable using observers with XEVE. We have implemented a prototype tool that purifies ESTEREL programs based on the algorithm.

2 ESTEREL Program Purification

In this section we describe how to transform a program that complies with certain requirements, into a pure program with an identical external observable behavior. We describe only the fundamentals of the algorithm[3].

We want to verify several safety properties of a given program using observers that monitor the system's behavior and emit special signals once one of these properties is violated. The observers are assembled in threads parallel to the main program. We would be able to verify that these signals are never emitted by compiling the observer-augmented program into BLIF format; however, we cannot do so as long as the program handles data other than pure signals.

The algorithm removes from the observed program all variables, sensors, and valued signals. Instead, we represent their values and statuses using pure signals. The value of a Boolean-valued object is represented by a single pure signal (presence represents *true*, while absence stands for *false*). As for numerical-valued objects whose values range over infinite domains, we partition their domains into non-overlapping intervals, which we hereby call *ranges*, in such way that two goals are achieved. The first one is being able to decide any condition containing an occurrence of some numerical valued object by knowing the range within which that object's value resides. For example, the condition `?s > 3.0` for a sensor s (`?s` is the current value of s) can be decided if the information whether $?s \in (-\infty, 3.0]$ or $?s \in (3.0, +\infty)$ is available. The second goal is maintaining relationships between dependent objects. For instance, for the assignment v `:=` a `*` `?`x `+` b, where v is a variable, x is a valued signal and a and b are literal constants, we can determine the range within which the v's value

[3] Due to lack of space we omit some details. The full description can be found in [8].

would reside following the assignment based on the range of x. E.g., for v `:= 2 * ?`x `+ 1` where `?`x is in (1,3], the value of v is in $(2 \cdot 1 + 1, 2 \cdot 3 + 1] = (3, 7]$. Note that actually we calculate the ranges for x given the partition for v.

The partitioning process starts from constant values assigned to the valued objects and Boolean data expressions in which they occur. It continues based on dependencies between valued objects, i.e., relations between valued objects formed when one object's value is assigned to another, for example, the variable assignment statement v `:= 3 *` u `+ 1` creates a dependency between variables v and u.

Using the calculated ranges we define pure signals that represent in the abstract system the values and statuses of valued objects from the concrete system. For example, an occurrence of an input signal X carrying a value x which resides in the *i-th* range calculated for X becomes an event in the abstract program in which the pure input signal `R`i`_`X is received from the environment.

The operations that handle data in the concrete program are replaced with pure signal operations having identical results (up to employing pure signals where valued objects are used in the original program).

Our method supports only programs that compile and execute without errors and that comply with several constraints. These constraints are required for maintaining relationships between valued objects and for replacing numerical and Boolean calculations with pure signal operations. Some are due to technical issues with aspects of our abstraction process conflicting with the synchrony hypothesis. One must remember that one way to allow automatic methods to check a non-trivial property is to reduce the power of the language or, equivalently, reduce the class of program verifiable using the method.

Our algorithm and the tool also make several additional assumptions about the program in order to simplify the solution. These assumptions usually concern ESTEREL syntactic sugaring instructions. However, they do not reduce the expressive power of ESTEREL: each assumption can be attained by replacing the original construct by a semantically-equivalent construct, and if the program has not been initially developed complying with these assumptions, there exist automatic procedures to simplify the original program. The assumptions are listed and explained in [8].

Combining the verification power of XEVE with the transformation of non-Pure ESTEREL programs into Pure ESTEREL programs, we can verify safety properties of a larger family of programs. Various robot behaviors appearing in [7] can be implemented in ESTEREL, processing pure signals or requiring sufficiently simple calculations, and thus our method can be applied to them.

In [8] we characterize the set of programs to which our method is applicable and present the application of our method to various control programs. For example:

– Bang-bang controllers[4] such as a temperature controller system turning on and off boilers to regulate the temperature in a chamber.
– Proportional controllers[5], e.g., a program regulating the speed of a vehicle by commanding the motor to accelerate in proportion to the difference between the target speed and the current speed.

Several examples of robot behaviors, demonstrating the application of the method to mobile robot programming following the *reactive control paradigm* are also provided in [8]. This paradigm, based on animal models of intelligence, decomposes the overall action of the robot by behavior, allowing handling multiple goals and multiple sensors, increasing robustness and extensibility [5]. By combining various behaviors and control algorithms, complicated control systems to which our method is applicable can be derived. The examples in [8] include original program code, abstract program code, observer code, simulations, and verification results.

3 Discussion

The program abstraction that we have implemented has an advantage over complete data abstraction performed when providing the `-soft` flag to the ESTEREL compiler, since it avoids adding behaviors not displayed by the original program. For example, consider the following code portion switching on and off an actuator based on a numerical input through a sensor `S`:

```
if ?S < 90.0 then emit On end if;
if ?S > 110.0 then emit Off end if
```

The ESTEREL compiler generates with the `-soft` flag a circuit in which both conditions can be satisfied at the same time, therefore it can emit both `On` and `Off` at the same reaction. However, no two range signals[6] representing `S` can occur simultaneously in the purified program when using our abstraction, hence both conditions cannot be satisfied at the same time.

By giving up completeness, the full potential of the technique developed is realized. Applying the technique to parts of the program that fulfill the constraints

[4] A bang-bang controller is a feedback controller that switches abruptly between two states [7]. The controller receives a measured quantity of interest, and outputs a certain value if that quantity is above a certain threshold, and a different value otherwise.

[5] A proportional controller is a closed-loop feedback controller whose control signal is proportional to the *error* – the difference between the *set point*, also known as the *reference* (the ideal point) and the measured quantity under control, i.e., the control signal is calculated by multiplying the *error signal* by a *gain* [7].

[6] In the abstraction we represent a sensor `S` using pure input signals, one for each range calculated for `S`: `R1_S`, `R2_S`, ..., `Rn_S`, which we call *range signals*. We declare an exclusion relation among the range signals such that the environment is incapable of providing more than one range signal at a time, and add a code segment that internally-emits the first range signal in every instant in which the other are absent, such that one is always present.

while letting the ESTEREL compiler remove the rest of the data when compiling the program into a circuit can produce a more precise over-approximation than total control-based abstraction. This is especially useful when the system consists of several sub-systems, some of which fulfilling the constraints while others not. An example of a system comprised of a PID controller[7] and a limit switch is provided in [8]. The limit switch shuts the process down by cutting off the PID controller's output once an undesired limit is reached. After the measured value drops back to the safe zone, the switch can be manually reset in order to reactivate the control system. An observer helps verifying the safety of the system by emitting a special signal whenever the program emits the output signal while the measured value is higher than some critical threshold. Not only the calculations performed by the PID controller are not supported by our technique, but also the program takes the set point, clock interval, and gains from constants defined in the host language. The abstraction performed by the ESTEREL compiler alone produces a circuit which XEVE reports to possibly emit the observer signal. However, using our technique to abstract the safety limit switch and the observer, which comply with the requirements of our techniques, and letting the compiler abstract the rest of the program creates a circuit which never emits the observer signal, as XEVE guarantees.

In future versions we intend to ease or completely overcome some of the constraints found in the current version and integrate our tool with other tools used for observer-based verification of ESTEREL programs.

References

1. G. Berry. The Esterel v5 Language Primer, Version v5_91. Centre de Mathématiques Appliquées – Ecole des Mines and INRIA, Sophia-Antipolis, France, 2000.
2. G. Berry et al. *The Esterel v5_91 System Manual*. Centre de Mathématiques Appliquées – Ecole des Mines de Paris / INRIA, Sophia-Antipolis, 2000.
3. A. Bouali. Xeve, an Esterel verification environment. In *CAV*, volume 1427 of *LNCS*, pages 500–504. Springer, 1998.
4. S. Das, D. L. Dill, and S. Park. Experience with predicate abstraction. In *CAV*, volume 1633 of *LNCS*, pages 160–171. Springer, 1999.
5. G. Dudek and M. Jenkin. *Computational Principles of Mobile Robotics*. Cambridge University Press, 2000.
6. N. Halbwachs. *Synchronous Programming of Reactive Systems*. Kluwer, 1993.
7. J. L. Jones and D. Roth. *Robot Programming: A Practical Guide to Behavior-Based Robotics*. McGraw-Hill, 2003.
8. N. Koblenc. Purification of Esterel Programs. Master's thesis. In preparation. Expected submission October 2014. The thesis draft and the prototype tool are available at http://www.cs.tau.ac.il/~tyshbe/NIR/nirThesisDraft.html.
9. I. Sommerville. *Software Engineering*. Addison-Wesley, 8 edition, 2007.
10. Esterel Technologies. The Esterel v7 Reference Manual, Version v7_30 – initial IEEE standardization proposal, 2005.

[7] A Proportional-Integral-Derivative PID controller is a controller whose control signal is the sum of three terms: one proportional to the error, one proportional to the derivative of the error, and one proportional to the integral of the error [7].