Formal Verification of an Executable LTL Model Checker with Partial Order Reduction

DOI:

[10.1007/s10817-017-9418-4](https://doi.org/10.1007/s10817-017-9418-4)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Brunner, J., & Lammich, P. (2018). Formal Verification of an Executable LTL Model Checker with Partial Order Reduction. *Journal of Automated Reasoning*, 60(1), 3-21. <https://doi.org/10.1007/s10817-017-9418-4>

Published in:

Journal of Automated Reasoning

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

Formal Verification of an Executable LTL Model Checker with Partial Order Reduction^{*}

Julian Brunner and Peter Lammich

Technische Universität München

Abstract. We present a formally verified and executable on-the-fly LTL model checker that uses ample set partial order reduction. The verification is done using the proof assistant Isabelle/HOL and covers everything from the abstract correctness proof down to the generated SML code. Building on Doron Peled’s paper “Combining Partial Order Reductions with On-the-Fly Model-Checking”, we formally prove abstract correctness of ample set partial order reduction. This theorem is independent of the actual reduction algorithm. We then verify a reduction algorithm for a simple but expressive fragment of PROMELA. We use static partial order reduction, which allows separating the partial order reduction and the model checking algorithms regarding both the correctness proof and the implementation. Thus, the CAVA model checker that we verified in previous work can be used as a back end with only minimal changes. Finally, we generate executable SML code using a stepwise refinement approach. We test our model checker on some examples, observing the effectiveness of the partial order reduction algorithm.

1 Introduction

Partial order reduction [25] is an important optimization for model checkers, enabling them to deal better with models involving concurrency. It allows the model checker to consider only a subset of all possible interleavings of concurrently executing operations by identifying equivalences between them. Unfortunately, partial order reduction is notoriously complex and can easily affect the correctness of the model checker. For instance, [25] describes a partial order reduction algorithm and claims that it can simply be used with on-the-fly nested depth-first search. It was found out later that this compromises correctness due to the reduction possibly differing between the inner and the outer search [8]. Moreover, while formalizing the algorithm in [25], we discovered that its correctness proof uses an invalid lemma (see section 2.2).

There is also the issue of implementation correctness, which is usually addressed via testing in the context of model checking algorithms. Since testing is necessarily incomplete, it may lead to incorrect implementations due to missed corner cases. Furthermore, when using models of realistic size, determining the correct outcome for a given test input requires the use of a model checker.

^{*} Research supported by DFG grant CAVA (Computer Aided Verification of Automata)

Thus, although in widespread use, neither the correctness of partial order reduction algorithms, nor the correctness of their implementations can be taken for granted. This is especially problematic since the trust in the correctness of a single model checker is used to justify the confidence in the correctness of the many models that it checks. In order to meet the very strict correctness requirements of model checking algorithms, we implement and formally verify a partial order reduction algorithm.

In previous work [5], we have presented the CAVA model checker, a fully verified and executable LTL model checker à la SPIN. The verification was done with the proof assistant Isabelle/HOL [24] and covers everything from the correctness of the algorithms down to the implementation. Due to its LCF-like architecture, Isabelle/HOL is more trustworthy than a large unverified implementation like SPIN (see section 3.1). This paper now adds the following contributions:

1. Formalization of a fragment of the modeling language PROMELA
2. Formalization of the static analysis required for partial order reduction
3. Formal abstract correctness proof for ample set partial order reduction
4. Verified implementation and integration into the CAVA model checker
5. Development of reusable libraries for automata and trace theory

This results in what we believe to be the first formally verified and executable implementation of a partial order reduction algorithm, addressing both of the issues mentioned earlier. The verification is carried out completely in Isabelle/HOL, such that the correctness of the model checker only depends on the correctness of Isabelle/HOL. This integration avoids logical gaps that may arise when manually composing the results of different verification tools. Most importantly, we now have a formally verified reference implementation that can deal with many formerly infeasible models, improving its usefulness for testing other model checkers.

To the best of our knowledge, there has been only one other attempt at formalizing partial order reduction [4]. However, it does not cover the reduction algorithm and is restricted to a specific fairness assumption (see section 2).

The rest of this paper is organized as follows. In section 2, we cover theoretical aspects of partial order reduction and elaborate on our choice of algorithm. In section 3, we report on our Isabelle/HOL formalization. In section 4, we compare the performance of our model checker to that of SPIN. Finally, in section 5, we give conclusions and future research directions.

2 Theory

Figure 1 illustrates the basics of partial order reduction. In regular model checking, the system automaton ‘ S ’ is derived from the system and used as input for the model checker together with the formula ‘ φ ’. The model checker then determines if the system automaton satisfies the property expressed by the formula ($\mathcal{L} S \subseteq \mathcal{L} \varphi$). When using partial order reduction, a reduction algorithm obtains a reduced system automaton ‘ R ’ from the system instead, which fulfills certain *reduction*

conditions. These conditions imply stuttering equivalence between the language of the system automaton and that of the reduced system automaton ($\mathcal{L} S \approx \mathcal{L} R$). Since properties expressed by next-free LTL formulae are stuttering-invariant [26], using the reduced system automaton instead of the system automaton when model checking yields the same result ($\mathcal{L} S \subseteq \mathcal{L} \varphi \iff \mathcal{L} R \subseteq \mathcal{L} \varphi$).

Fig. 1: Partial Order Reduction Overview. A reduction algorithm obtains the reduced system automaton ‘ R ’, which is then used as an input of the model checker instead of the system automaton ‘ S ’. The reduction algorithm guarantees that the reduced system automaton fulfills certain reduction conditions, from which one can prove stuttering equivalence between the two languages. This implies that the result of the model checker is not affected by the reduction.

Note that this is a very abstract description of partial order reduction. In actual implementations, the reduced system automaton may be represented implicitly, and the reduction algorithm may be merged with the model checking algorithm. However, this view allows us to identify the three major tasks involved in developing a verified implementation of partial order reduction:

1. Reduction algorithm correctness: The automaton produced by the reduction algorithm fulfills the reduction conditions.
2. Abstract correctness: If an automaton fulfills the reduction conditions, its language is stuttering equivalent to that of the system automaton.
3. Implementation and verification of the reduction algorithm.

Unlike our formalization, [4] only covers the second task. This means there is no input language, no static analysis, no reduction algorithm, no implementation, and no executable model checker. Furthermore, it only covers the case where

a certain fairness assumption is met, which simplifies the abstract correctness proof. In absence of other formalization attempts, we believe that our work is a significant contribution over the existing body of research.

2.1 Reduction Conditions

Both the reduction algorithm and the abstract correctness are built around the reduction conditions, making them the main object of interest when dealing with partial order reduction. We chose to implement an algorithm based on the ample set method and chose the reduction conditions accordingly. Let ‘en q ’ be the set of enabled actions at state ‘ q ’ of the system automaton (*enabled set*). Let ‘ren q ’ be the set of enabled actions at state ‘ q ’ of the reduced system automaton (*ample set*). Let ‘ex $a q$ ’ be the successor of state ‘ q ’ after executing action ‘ a ’ (‘ex’ is called *execution function*). This way, ‘(en, ex)’ represents the system automaton, while ‘(ren, ex)’ represents the reduced system automaton. The set of finite words executable at state ‘ q ’ of the system automaton ‘words q ’ is defined in terms of ‘en’ and ‘ex’. For a more detailed description of the system definitions, see section 3.4. With these prerequisites, we define the following reduction conditions:

subset	$\forall q. \text{ren } q \subseteq \text{en } q$
nonempty	$\forall q. \text{ren } q \subset \text{en } q \implies \text{ren } q \neq \{\}$
independent	$\exists \text{independence relation } I. \forall q w. \text{ren } q \subset \text{en } q \implies$ $w \in \text{words } q \implies \text{ren } q \cap \text{set } w = \{\} \implies I(\text{ren } q)(\text{set } w)$
wellfounded	$\exists \text{well-founded relation } R. \forall q a. \text{ren } q \subset \text{en } q \implies$ $a \in \text{ren } q \implies R(\text{ex } a q) q$
invisible	$\forall q. \text{ren } q \subset \text{en } q \implies \text{ren } q \subseteq \text{invisible}$

Condition **subset** states that the reduced system automaton is a subautomaton of the system automaton and is usually not stated explicitly in the literature. Condition **nonempty** states that the reduction algorithm must not omit all of the actions at any state. Condition **independent** requires that all the actions that are executed after reaching some state but before an action from the ample set at this state are *independent* of all the actions in this ample set. Condition **wellfounded** requires that every cycle in the system automaton contains at least one state where no reduction is performed. Condition **invisible** states that when a proper reduction takes place, the ample set cannot contain any actions that are *visible* to the formula. Conditions **nonempty**, **independent**, and **wellfounded** correspond to conditions C0, C1, and C2 in [4, pages 268, 269], while condition **invisible** corresponds to condition C3’ in [25, page 50]. Note that even though the reduction conditions are similar, our formalization is not based on [4].

2.2 Reduction Algorithm

These conditions are very abstract, so there are still many choices to be made with respect to the actual reduction algorithm. We originally planned to verify dynamic partial order reduction with on-the-fly model checking [25], but soon encountered

difficulties. Dynamic partial order reduction detects cycles during the emptiness check in order to ensure condition **wellfounded**. This tight integration with the emptiness check has led to bugs in the past [8]. When used with on-the-fly model checking, this integration also extends to the product construction, effectively turning the whole system into one monolithic algorithm. It also introduces a mismatch since an algorithm that conceptually works on a system automaton is now used with a product automaton, requiring complicated reasoning. And indeed, during our effort of formalizing the proof given in [25], we discovered a counterexample for one of the lemmata used in this proof. This counterexample is based on the fact that, when exploring the product automaton, different instances of the system automaton appearing in the product automaton may be reduced differently. A more detailed description can be found in [3, section 8.4]. Note that this, while refuting the lemma, does not necessarily invalidate the correctness theorem, only the proof thereof. However, despite investing a significant amount of time, we were unable to find an alternative proof as it seems that the reasoning required is more complex than anticipated in the original paper.

We chose to implement a static partial order reduction [9] algorithm instead, which avoids these problems of the dynamic approach. It ensures condition **wellfounded** by performing some static analysis initially, identifying a set of *sticky* edges which breaks every cycle in the control flow graph. Static partial order reduction is much more modular, making it possible to verify the reduction algorithm independently of the product construction and the emptiness check. This way, we were able to simply add the reduction algorithm as a preprocessing step to the existing CAVA model checker, enabling reuse of existing optimizations.

The reduction algorithm itself is similar to the one used in SPIN [7]. The basic idea is to take the set of enabled actions of each process in the state as a candidate for an ample set. For each candidate, an over-approximation of the reduction conditions is tested. If no candidate satisfies the conditions, the state is fully expanded, that is, no reduction is performed.

For instance, our approximation checks that, in order to be used as an ample set, the actions of a process must be independent of all actions of other processes. Moreover, it is checked that no additional action of this process can be enabled as a consequence of executing actions of other processes. Thus, only independent actions of other processes can be executed before an action of the ample set, which implies condition **independent**.

3 Formalization

Our formalization contains all three of the tasks outlined in section 2. The implementation was integrated into the CAVA model checker, which was published previously [5, 6]. Since then, various features have been added to this model checker. For instance, it now supports using PROMELA as an input language [22]. Furthermore, the library for automata has been updated [13] and a new framework for depth-first search algorithms has been formalized [16]. Also, an alternative algorithm for deciding language emptiness of Büchi automata based

on Gabow’s strongly-connected components algorithm has been implemented [14]. However, the focus of this paper is on the implementation and verification of the partial order reduction algorithm.

In this section, we give some technical background regarding the tools that were used as well as a high-level overview of the formalization. We also describe certain noteworthy aspects of the formalization in isolated detail. The full formalization is available at https://cava.in.tum.de/CAVA_POR.

3.1 Isabelle/HOL

Isabelle/HOL [24, 23] is a proof assistant based on Higher-Order Logic (HOL), which can be thought of as a combination of functional programming and logic. Formalizations done in Isabelle/HOL are trustworthy for two reasons. Firstly, Isabelle’s LCF architecture guarantees that all proofs are checked using a very small logical core which is rarely modified but tested extensively over time. This reduces the trusted code base to a minimum. Secondly, bugs in the core rarely lead to accidentally proving false propositions. Bugs that have large effects are easily caught, while the limited applicability of bugs with small effects is unlikely to coincide with a logical mistake in the large-scale structure of the proof.

Isabelle/HOL notation resembles standard mathematical notation with just a few differences. For instance, as in functional programming, functions are usually curried in HOL. This means that instead of ‘ $f :: A \times B \rightarrow C$ ’ with application syntax ‘ $f(x, y)$ ’, we have ‘ $f :: A \rightarrow B \rightarrow C$ ’ with application syntax ‘ $f\ x\ y$ ’.

3.2 Refinement Framework

When developing formally verified algorithms, there is a trade-off between the efficiency of the algorithm and the efficiency of the proof: For complex algorithms, a direct proof of an efficient implementation tends to get unmanageable, as implementation details obfuscate the main ideas of the proof. A standard approach to this problem is stepwise refinement [1], which modularizes the correctness proof: One starts with an abstract version of the algorithm and then refines it in correctness preserving steps to the concrete, efficient version. A refinement step may reduce the nondeterminism of a program, replace abstract mathematical specifications by concrete algorithms, and replace abstract datatypes by their implementations. For example, selection of an arbitrary element from a set may be refined to getting the head of a list. This approach separates the correctness proof of the algorithm, which focuses on the main algorithmic ideas, from the correctness proof of the implementation, where the proof of each refinement step focuses on a specific implementation detail, not caring about the overall correctness property.

In Isabelle/HOL, stepwise refinement is supported by the Refinement Framework [11, 17] and the Isabelle Collection Framework [10, 15]. The former framework implements a refinement calculus [1] based on a nondeterminism monad [27], and the latter provides a library of verified efficient data structures. Both frameworks

come with tool support to simplify their usage for algorithm development and to automate canonical tasks such as verification condition generation.

3.3 Basics

The most basic concept needed for nearly all parts of the formalization is that of *sequences*. With HOL being very similar to functional programming languages like SML or Haskell, the standard library already includes extensive support for *finite sequences* via the type ‘ α list = Nil | Cons α (α list)’ . For *infinite sequences*, the type ‘ α word’ is used, which is simply a type synonym for ‘ $\mathbb{N} \rightarrow \alpha$ ’.

We also use the library Coinductive [18] which formalizes lazy lists using codatatypes [2]. It provides the type ‘ α llist’, which models both finite and infinite sequences. This is useful for selecting subsequences of infinite lists that can be either finite or infinite. Reasoning about selections and indices of lazy lists required us to significantly extend the library Coinductive.

Another important component needed for partial order reduction is stuttering equivalence and the proof that next-free LTL formulae can only express stuttering-invariant properties. The library Stuttering Equivalence [20] is used for both.

3.4 Systems

Model checkers usually represent systems using the type ‘(state $\times$ state) set’. Reasoning about partial order reduction requires transitions to be labeled with actions, suggesting the type ‘(state $\times$ action $\times$ state) set’. However, this type allows multiple successor states to be reached given a state and an action, making the type a bad fit for the deterministic action model of partial order reduction. This leads to unnecessary wellformedness conditions, inaccessible successor states, and overspecified path predicates. We thus chose the following representation of the system automaton which was already referred to in section 2.1:

$$\text{en} :: \text{state} \rightarrow \text{action set} \quad (1a)$$

$$\text{ex} :: \text{action} \rightarrow \text{state} \rightarrow \text{state} \quad (1b)$$

$$\text{init} :: \text{state set} \quad (1c)$$

Here, ‘en’ is the set of enabled actions at a state (*enabled set*), ‘ex’ is the function that, given an action, maps each state to its successor state (*execution function*), and ‘init’ is the *set of initial states*.

This representation allows paths to be introduced in a straightforward way via the inductively defined set ‘words :: state $\rightarrow$ action list set’:

$$[] \in \text{words } p \quad (2a)$$

$$a \in \text{en } p \implies w \in \text{words } (\text{ex } a \ p) \implies a \# w \in \text{words } p \quad (2b)$$

Inductive definitions in Isabelle/HOL specify the smallest sets that satisfy the given rules. Equivalently, they specify the sets containing those elements whose membership can be derived using the given rules. These rules can be declared

as safe introduction rules, so that whenever Isabelle/HOL encounters proof obligations of the form ‘ $[] \in \text{words } p$ ’ or ‘ $a \# w \in \text{words } p$ ’, it can automatically split them into simpler goals or discharge them completely.

We prove an additional rule for the append operator on lists:

$$u \in \text{words } p \implies v \in \text{words } (\text{fold ex } u \text{ } p) \implies u @ v \in \text{words } p \quad (3)$$

Note how ‘fold’ lifts the execution function ‘ $\text{ex} :: \text{action} \rightarrow \text{state} \rightarrow \text{state}$ ’ from single actions to sequences of actions ‘ $\text{fold ex} :: \text{action list} \rightarrow \text{state} \rightarrow \text{state}$ ’. Also note how this rule generalizes rule 2b.

Together, rules 2a, 2b, and 3 form a set of introduction rules that break down most goals automatically. For instance, the goal ‘ $u @ a \# v \in \text{words } p$ ’ gets transformed into three subgoals:

$$u \in \text{words } p \quad (4a)$$

$$a \in \text{en } (\text{fold ex } u \text{ } p) \quad (4b)$$

$$v \in \text{words } (\text{ex } a \text{ } (\text{fold ex } u \text{ } p)) \quad (4c)$$

This automates proofs significantly, in some cases shortening proofs comprised of 50 to 100 lines to one-liners. We have proven many more rules about this system formalization, making it a useful addition to the CAVA automata library.

3.5 Trace Theory

Partial order reduction introduces the concept of *independent* actions, which can be executed in any order without changing the result or enabling or disabling each other. Trace theory [19] lifts this notion of commutable items to that of *equivalent* sequences, which is needed in the abstract correctness proof.

Finite sequences are equivalent if they differ by a finite number of commutations of independent actions. This concept is then extended to infinite sequences [25, page 41]. This definition by case distinction makes lazy lists difficult to use, so we decided to work with separate types and definitions for finite and infinite sequences.

Formalizing the necessary parts of trace theory took significant effort due to the large number of theorems. There are also some theorems that look simple but are difficult to prove, for instance:

$$w_1 \equiv_I w_2 \iff u @ w_1 @ v \equiv_I u @ w_2 @ v \quad (5)$$

The left to right direction can be proven via rule induction on the transitive structure of ‘ $\equiv_I$ ’. Doing the same for the right to left direction results in an unprovable induction step. It was necessary to prove the following lemmata:

$$w_1 \equiv_I w_2 \implies \text{remove1 } c \text{ } w_1 \equiv_I \text{remove1 } c \text{ } w_2 \quad (6a)$$

$$u @ w_1 \equiv_I u @ w_2 \implies w_1 \equiv_I w_2 \quad (6b)$$

$$w_1 \equiv_I w_2 \implies \text{rev } w_1 \equiv_I \text{rev } w_2 \quad (6c)$$

Here, ‘remove1 c w ’ removes the first occurrence of ‘ c ’ from the sequence ‘ w ’, and ‘rev w ’ reverses the sequence ‘ w ’. Lemma 6a uses ‘remove1’ to avoid the fact that rule induction does not work with modified assumptions. We use lemma 6a to prove lemma 6b via reverse induction on the sequence ‘ u ’. Lemma 6c is proven via rule induction and with lemma 6b, it completes the proof of theorem 5.

We also had to define some concepts specific to partial order reduction. For instance, the predicate specifying that the first occurrence of a symbol in a sequence is independent of all symbols before it. In the end, the formalization of the relevant aspects of trace theory required about as much proof text as the formalization of the abstract correctness proof itself.

3.6 Abstract Correctness

Assume that ‘ S ’ is a system automaton and ‘ R ’ is a reduced system automaton such that the reduction conditions introduced in section 2.1 hold. Then, the abstract correctness theorem states that the languages of ‘ S ’ and ‘ R ’ are stuttering equivalent:

$$\mathcal{L} S \approx \mathcal{L} R \quad (7)$$

The proof of this theorem required about 1000 lines of formal proof text and dozens of lemmata. Its structure is similar to that of the informal proof [25] and we will thus not repeat it here.

However, we present the formalization of a lemma [25, Theorem 3.11] in detail and highlight the differences between the formal and the informal proof:

```

lemma reduction_word:
  assumes "q ∈ reachable" "v ∈ wordsS q"
  obtains u w
  where
    "w ∈ wordsR q"
    "v ≡I u" "u ≼I w"
    "lproject visible (inf_llist u) = lproject visible (inf_llist w)"

```

Note that we do not present the formal definitions of all the constants used in this theorem. Informally, the theorem states that, given an infinite sequence ‘ v ’ in the system automaton, it is possible to find a corresponding sequence ‘ w ’ in the reduced system automaton. The theorem also implies the existence of an intermediate sequence ‘ u ’, which is needed since ‘ w ’ may contain actions that are not in ‘ v ’.

The proof consists of two parts. In the first part, we construct an arbitrarily long but finite sequence in the reduced system automaton by transcribing longer and longer prefixes of the infinite sequence in the system automaton. In order to do so, we inductively define a predicate that describes a valid state during this construction process where a prefix of the sequence in the system automaton has already been processed. This predicate specifies that the state of the construction

where both the sequence in the system automaton and the one in the reduced system automaton are empty is valid. It also specifies how one can extend a valid construction state by adding a step in the system automaton and a sequence of corresponding steps in the reduced system automaton. At each point of the construction, we can then prove that some invariants hold and that the construction can be extended. Proving these invariants and the extension property required a lot of effort as the informal proof only provided a rough sketch of the argument. The formal proof constitutes both a certificate of the theorem’s correctness as well as a detailed documentation of the reasoning used to prove it.

The second part of the proof consists of using the first part to show that there exists an infinite sequence with the required properties in the reduced system automaton. While this step is almost completely skipped in the informal proof, the formal one forces us to consider it rigorously. For instance, the first part supplies a theorem which guarantees that for any number of steps that were already taken, another step can be taken, extending the sequence in the process. Intuitively, such a theorem can be applied “infinitely often” to obtain an infinite sequence, but this is not logically sound. Performing a step like this in a formal proof requires precise reasoning and in our case the use of Hilbert’s epsilon operator. We believe that this is not a flaw of formal logic or the particular instance we are using. Instead, we think that situations like this point to areas where it became customary to use sloppy reasoning in informal proofs, possibly leading to mistakes or overlooked side conditions. For instance, it is often not made clear in which way variables depend on each other or what guarantees that an infinite sequence can actually be constructed from a set of finite sequences. Formal proofs point out required side conditions like the fact that the infinite concatenation of these finite sequences needs to be infinite. It also brought attention to the fact that many concepts need to be defined on both finite and infinite sequences and that they need to correspond to each other in a specific way.

As mentioned in sections 3.3, 3.4, and 3.5, a large amount of foundational work was required in order to formally prove the abstract correctness theorem.

3.7 The SM Language

In order to implement an executable reduction algorithm, we require a concrete modeling language. We use a simple fragment of PROMELA that is expressive enough to model interesting examples. We call this fragment the *SM language*.

A program in this language consists of a set of processes, each of which is described using a guarded command language. Each process has a set of local variables and communication between processes is modeled via global variables. A configuration of the system consists of a valuation of the global variables and a list of process configurations, where a process configuration consists of a command and a valuation of the local variables. The main PROMELA feature not supported by SM is channels, which can be emulated by global variables.

We specify a structural operational semantics that establishes a control flow graph where the nodes are commands and the edges are labeled with *local actions*. A local action can be a guarded assignment, a test, or the skip action. Each local

action is assigned an enabledness check and an effect function on the local and global variables.

The system semantics describes a step relation between configurations by nondeterministically picking a process from a configuration, following an edge in the control flow graph from the process' command that is labeled with an enabled local action, and applying the effect of the local action to the local and global state. To ensure that all runs of the system are infinite, we apply a stuttering extension, that is, if there is no process with an enabled action, the system may take a step that does not change the configuration.

Since we want to use the SM language in an LTL model checker, we need to define *atomic propositions* and their connection to the system states. In our case, atomic propositions are simply expressions in the SM language that contain only global variables. Then, we define the *interpretation function* to map each state to the set of expressions that evaluate to true in this state.

We define the language of a program as the set of infinite sequences of sets of atomic propositions that correspond to infinite runs of the program:

$$\mathcal{L} :: \text{program} \rightarrow \text{exp set word set} \quad (8)$$

We define a *global action* to consist of a process id and a control flow graph edge. The process id is the position of the associated process in the list of all processes. A global action is enabled if the associated process exists, the control flow graph edge is consistent with the current command of the associated process, and the corresponding local action is enabled. Execution of a global action transforms the state of the associated process and the global variables according to the corresponding local action.

3.8 Reduction Algorithm

Next, we define a function that selects an ample set for a configuration. Similar to SPIN, candidates for ample sets are the sets of enabled actions of each process. We make a rather crude approximation and allow a nonempty set of enabled actions of a process as an ample set, if (1) there is no statically enabled action of the process that reads or writes global variables, and (2) none of the enabled actions corresponds to a sticky edge in the control flow graph. Here, (1) is a simple way of guaranteeing condition **independent** (see section 2.1), and (2) is the condition imposed by static partial order reduction (see section 2.2).

We implemented and verified an algorithm based on depth-first search which computes the set of sticky edges before the actual model checking phase. This algorithm starts with the set of edges labeled with actions containing global variables and extends it to a feedback arc set on the control flow graphs of the processes. For this task, we used the Depth-First Search Framework [16], which simplifies the implementation and verification of efficient DFS-based algorithms.

We define the reduced system automaton based on this ample function and prove that all of the reduction conditions from section 2.1 are fulfilled. This allows us to invoke the abstract correctness theorem to obtain stuttering equivalence

between the language of the system automaton and that of the reduced system automaton. Together with the assumption that the formula is next-free, this implies that using the reduced system automaton for model checking instead of the system automaton does not change the result.

3.9 Integration

We refine the ample function, the execution function, and the interpretation function to efficiently executable implementations. Among other steps, this includes compilation of the model to a more efficient representation. Finally, instantiating the generic infrastructure of the CAVA model checker yields an executable LTL model checker ‘cava’ which uses the reduced system automaton. Obtaining the main theorem of our development is then merely a matter of combining the correctness theorem of the CAVA model checker with that of abstract partial order reduction:

$$\text{case cava } S \varphi \text{ of SAT} \Rightarrow \mathcal{L} S \subseteq \mathcal{L} \varphi \mid \text{UNSAT} \Rightarrow \mathcal{L} S \not\subseteq \mathcal{L} \varphi \quad (9)$$

This theorem states that the function ‘cava’ decides whether or not the sequences of atomic propositions admitted by runs of the program satisfy the LTL formula. The meaning of this statement only depends on the abstract semantics of the SM language (term ‘ $\mathcal{L} S$ ’) and the abstract semantics of LTL formulae (term ‘ $\mathcal{L} \varphi$ ’). All other parts of the formalization, including partial order reduction, LTL model checking, and refinement towards efficiently executable definitions, are covered by this machine-checked correctness theorem. Note that we also formalized a version of the model checker that provides a counterexample in case the program does not satisfy the formula.

Finally, Isabelle/HOL can generate Standard ML code from the definition of the function ‘cava’. This code then constitutes a formally verified and executable LTL model checker. A snapshot of this formalization can be found at https://cava.in.tum.de/CAVA_POR.

We conclude with some statistics about the formalization, which took about 15 man-months and resulted in about 13k lines of theory text being added to the model checker. This includes both definitions and proofs and splits up into 6k lines for abstract partial order reduction and 7k lines for the SM language and the associated program analysis. The size of the whole codebase of the model checker and its libraries is about 140k lines of theory text.

4 Evaluation

We perform some basic sanity checks using two systems that admit no reduction and complete sequentialization, respectively. As a practical example, we implement a distributed mutual exclusion algorithm called MULOLOG [21] using the supported PROMELA fragment. The property used for testing states that at most one process can be in the critical section at any point in time. We perform model checking

Example	Processes	States SPIN	States SPIN*	States CAVA	States CAVA*
MULOG	1	27	27	52	52
MULOG	2	2,674	2,004	5,538	4,284
MULOG	3	2,376,180	1,171,578	5,205,376	2,779,218

Fig. 2: Reduction effectiveness. Shown are the number of states that were explored during model checking using both the CAVA and the SPIN model checkers. The starred variants indicate where partial order reduction was used.

using both the CAVA and the SPIN model checkers, both with and without partial order reduction. Figure 2 shows the reduction effectiveness for this algorithm.

Both the CAVA and the SPIN model checker show a significant reduction in the number of states. The reduction factors are comparable (roughly 1.3 for two processes and roughly 2 for three processes). The SPIN model checker explores fewer states in total (roughly factor 2) and has shorter execution times (roughly factor 400) than the CAVA model checker.

We would like to emphasize that in this paper, it is not our goal to compete with SPIN in absolute terms. Instead, our focus is on providing a verified and executable reference implementation of partial order reduction. The SPIN model checker employs various other optimizations and compilation to C code, while the CAVA model checker interprets the semantics of the modeling language. Thus, little insight can be gained by directly comparing execution time and memory consumption. Incorporating these optimizations is orthogonal to partial order reduction and we consider this subject of further research. Due to the modular architecture of the CAVA model checker, doing so will not make this contribution obsolete. At this point, it will also be possible to perform a more comprehensive evaluation with multiple example algorithms.

5 Conclusion

Formal verification is sometimes downplayed as “careful documentation of known theorems” or “filling in obvious details in proofs”. In practice, formal verification usually involves extensive modeling as well as abstraction, generalization, and simplification of the theory. What may seem like trivial completion of the informal proof often involves bridging large gaps and proving omitted corner cases.

In this project, we discovered an issue with the correctness proof given in [25] (see section 2.2). This demonstrates both the need for and the usefulness of formal verification. More importantly, we developed a formally verified and executable LTL model checker with partial order reduction. As the verification is machine-checked and covers everything from the abstract algorithm to the generated SML code, this is a very strong correctness guarantee. Our model checker is fast enough to serve as a reference implementation for other model checkers on models of realistic size. This constitutes a much-needed source of trust given the widespread use of partial order reduction together with its history of issues. The formalization can further serve as a detailed description of the

theory of partial order reduction and its correctness proof, which is useful since nontrivial gaps were bridged in the proof. We also developed a significant amount of foundational theories that can be reused in other projects dealing with similar concepts. Finally, our work demonstrates that large systems can now be verified using proof assistants via modularization and reuse of existing theories.

Future work consists of extending the SM language to make it more practical, with the ultimate goal of supporting most or all of the features of PROMELA. It is also possible to find smaller sets of sticky actions by incorporating heuristics about variable increments/decrements [9]. Another way to improve reduction consists of using additional static analysis to find larger independence relations. Finally, there is still room for improvement concerning the implementation, especially via the use of imperative data structures [12].

References

- [1] Ralph-Johan Back and Joakim von Wright. *Refinement Calculus. A Systematic Introduction*. Graduate Texts in Computer Science. Springer, 1998.
- [2] Jasmin Christian Blanchette, Johannes Hölzl, Andreas Lochbihler, Lorenz Panny, Andrei Popescu, and Dmitriy Traytel. “Truly Modular (Co)data-types for Isabelle/HOL”. In: *ITP*. Vol. 8558. LNCS. Springer, 2014, pp. 93–110.
- [3] Julian Brunner. “Implementation and Verification of Partial Order Reduction for On-The-Fly Model Checking”. MA thesis. Technische Universität München, July 15, 2014. 83 pp. URL: <http://www21.in.tum.de/~brunnerj/documents/ivporotfmc.pdf>.
- [4] Ching-Tsun Chou and Doron Peled. “Formal Verification of a Partial-Order Reduction Technique for Model Checking”. In: *TACAS*. Vol. 1055. LNCS. Springer, 1996, pp. 241–257.
- [5] Javier Esparza, Peter Lammich, René Neumann, Tobias Nipkow, Alexander Schimpf, and Jan-Georg Smaus. “A Fully Verified Executable LTL Model Checker”. In: *CAV*. Vol. 8044. LNCS. Springer, 2013, pp. 463–478.
- [6] Javier Esparza, Peter Lammich, René Neumann, Tobias Nipkow, Alexander Schimpf, and Jan-Georg Smaus. “A Fully Verified Executable LTL Model Checker”. In: *Archive of Formal Proofs* (May 2014). Formal proof development. URL: http://afp.sf.net/entries/CAVA_LTL_Modelchecker.shtml.
- [7] Gerard J. Holzmann. *The SPIN Model Checker. Primer and Reference Manual*. Addison-Wesley Professional, Sept. 2003.
- [8] Gerard J. Holzmann, Doron Peled, and Mihalis Yannakakis. “On Nested Depth First Search”. In: *SPIN Workshop*. Vol. 32. 1996, pp. 81–89.
- [9] Robert Kurshan, Vladimir Levin, Marius Minea, Doron Peled, and Hüsnü Yenigün. “Static Partial Order Reduction”. In: *TACAS*. Vol. 1384. LNCS. Springer, 1998, pp. 345–357.

- [10] Peter Lammich. “Collections Framework”. In: *Archive of Formal Proofs* (Nov. 2009). Formal proof development. URL: <http://afp.sf.net/entries/Collections.shtml>.
- [11] Peter Lammich. “Refinement for Monadic Programs”. In: *Archive of Formal Proofs* (Jan. 2012). Formal proof development. URL: http://afp.sf.net/entries/Refine_Monadic.shtml.
- [12] Peter Lammich. “Refinement to Imperative/HOL”. In: *ITP*. Vol. 9236. LNCS. Springer, 2015, pp. 253–269.
- [13] Peter Lammich. “The CAVA Automata Library”. In: *Archive of Formal Proofs* (May 2014). Formal proof development. URL: http://afp.sf.net/entries/CAVA_Automata.shtml.
- [14] Peter Lammich. “Verified Efficient Implementation of Gabow’s Strongly Connected Component Algorithm”. In: *ITP*. Vol. 8558. LNCS. Springer, 2014, pp. 325–340.
- [15] Peter Lammich and Andreas Lochbihler. “The Isabelle Collections Framework”. In: *ITP*. Vol. 6172. LNCS. Springer, 2010, pp. 339–354.
- [16] Peter Lammich and René Neumann. “A Framework for Verifying Depth-First Search Algorithms”. In: *CPP*. ACM, Jan. 13, 2015, pp. 137–146.
- [17] Peter Lammich and Thomas Tuerk. “Applying Data Refinement for Monadic Programs to Hopcroft’s Algorithm”. In: *ITP*. Vol. 7406. LNCS. Springer, 2012, pp. 166–182.
- [18] Andreas Lochbihler. “Coinductive”. In: *Archive of Formal Proofs* (Feb. 2010). Formal proof development. URL: <http://afp.sf.net/entries/Coinductive.shtml>.
- [19] Antoni Mazurkiewicz. “Trace Theory”. In: *Advances in Petri Nets, Part II*. Vol. 255. LNCS. Springer, 1987, pp. 278–324.
- [20] Stephan Merz. “Stuttering Equivalence”. In: *Archive of Formal Proofs* (May 2012). Formal proof development. URL: http://afp.sf.net/entries/Stuttering_Equivalence.shtml.
- [21] Mohamed Naimi, Michel Trehel, and André Arnold. “A Log (N) Distributed Mutual Exclusion Algorithm Based on Path Reversal”. In: *Journal of Parallel and Distributed Computing* 34.1 (1996), pp. 1–13.
- [22] René Neumann. “Using Promela in a Fully Verified Executable LTL Model Checker”. In: *VSTTE*. LNCS. Springer, 2014, pp. 105–114.
- [23] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL. A Proof Assistant for Higher-Order Logic*. Vol. 2283. LNCS. Springer, 2002.
- [24] Larry Paulson, Tobias Nipkow, and Makarius Wenzel. *Isabelle*. 2014. URL: <http://isabelle.in.tum.de>.
- [25] Doron Peled. “Combining Partial Order Reductions with On-the-Fly Model-Checking”. In: *Formal Methods in System Design* 8.1 (1996), pp. 39–64.
- [26] Doron Peled and Thomas Wilke. “Stutter-Invariant Temporal Properties are Expressible Without the Next-Time Operator”. In: *Information Processing Letters* 63.5 (1997), pp. 243–246.
- [27] Philip Wadler. “Comprehending Monads”. In: *Mathematical Structures in Computer Science* 2 (04 Dec. 1992), pp. 461–493.