

Symbolic Reachability Analysis of B through ProB and LTSmin

Jens Bendisposto, Philipp Körner, Michael Leuschel,
Jeroen Meijer, Jaco van de Pol,
Helen Treharne, Jordan Whitefield

This is a pre-print version archived by P. Körner at
<https://pkoerner.github.io/pages-output/publications/>.

The final publication is available at Springer via
https://dx.doi.org/10.1007/978-3-319-33693-0_18.

Symbolic Reachability Analysis of B through PROB and LTSMIN

Jens Bendisposto¹, Philipp Körner¹, Michael Leuschel¹, Jeroen Meijer^{2*}, Jaco van de Pol², Helen Treharne³, and Jorden Whitefield^{3†}

¹ Institut für Informatik, Heinrich Heine University Düsseldorf, Germany
{bendisposto@cs., p.koerner@, leuschel@cs.}uni-duesseldorf.de

² Formal Methods and Tools, University of Twente, The Netherlands
{j.j.g.meijer, j.c.vandepol}@utwente.nl

³ Department of Computer Science, University of Surrey, United Kingdom
{h.treharne, j.whitefield@surrey.ac.uk

Abstract. We present a symbolic reachability analysis approach for B that can provide a significant speedup over traditional explicit state model checking. The symbolic analysis is implemented by linking PROB to LTSMIN, a high-performance language independent model checker. The link is achieved via LTSMIN's PINS interface, allowing PROB to benefit from LTSMIN's analysis algorithms, while only writing a few hundred lines of glue-code, along with a bridge between PROB and C using $\text{\textcircled{O}MQ}$. PROB supports model checking of several formal specification languages such as B, Event-B, Z and TLA^+ . Our experiments are based on a wide variety of B-Method and Event-B models to demonstrate the efficiency of the new link. Among the tested categories are state space generation and deadlock detection; but action detection and invariant checking are also feasible in principle. In many cases we observe speedups of several orders of magnitude. We also compare the results with other approaches for improving model checking, such as partial order reduction or symmetry reduction. We thus provide a new scalable, symbolic analysis algorithm for the B-Method and Event-B, along with a platform to integrate other model checking improvements via LTSMIN in the future.

Keywords: B-Method, Event-B, PROB, LTSMIN, symbolic reachability

1 Introduction

In this paper we describe the process, technique and design decisions we made for integrating the two tooling sets: LTSMIN and PROB. Bicarregui et al. suggested, in a review of projects which applied formal methods [5], that providing useable tools remained a challenge. Recent use of the PROB tool in a rail system case study [15], where model checking large industrial sized complex specifications was performed, illustrated that there continues to be limitations with the tooling.

*Supported by STW SUMBAT grant: 13859

†Partly supported by EPSRC grant: EP/M506655/1




Fig. 1: Modular PINS architecture of LTSMIN [16]

Model checking $\text{CSP}\parallel\text{B}$ [32] specifications in PROB was the original motivator for this research, and based on a promising initial exploration [34], this paper defines a systematic integration of the two tooling sets.

LTSMIN is a high-performance language-independent model checker that allows numerous modelling language front-ends to be connected to various analysis algorithms, through a common interface, as shown in Figure 1. It offers a wide spectrum of parallel and symbolic algorithms to deal with the state space explosion of different verification problems. This connecting interface is called the PARTITIONED INTERFACE to the Next-State function (PINS), the basis of which consists of a state-vector definition, an initial state, a partitioned successor function (NEXTSTATE), and labelling functions [16]. It is through PINS that we have been able to leverage the PROB tool, therefore allowing us to take advantage of LTSMIN’s algorithmic back-ends. In this paper we focus on the new PROB language front-end, the grouping of transitions, and the symbolic back-end. In Section 5 we also briefly discuss state variable orders.

PROB [20] is an animator and model checker for many different formal languages [29], including the classical B-Method [2], Event-B [1], CSP, $\text{CSP}\parallel\text{B}$, Z and TLA^+ . PROB can perform automatic or step by step animation of B machines, and can be used to systematically verify the behaviour of machines. The verification can identify states which do not meet the invariants, do not satisfy assertions or that deadlock. At the heart of PROB is a constraint solver, which enables the tool to animate and model check high-level specifications. The built-in model checker is a straightforward, explicit state model checker (albeit augmented with various features such as symmetry reduction [21] or partial order reduction [10]). The explicit state model checker TLC can also be used as a backend [11].

The purpose of this paper is to make use of the advanced features of the LTSMIN model checker, such as symbolic reachability analysis, by linking the PROB state exploration engine with LTSMIN. This is achieved through a C programming interface [4] within the PROB tool, allowing the representation of a state to be compatible for LTSMIN’s consumption. In this paper the integration focuses on what is required in order to perform symbolic reachability analysis of

B-Method and Event-B specifications. The contribution of this research is a new tool integration, which can be used as a platform for further extensions.

The paper is structured as follows: Section 2 presents an overview of the B-Method, a running example and an illustration of definitions of transition systems used by LTSMIN. Section 3 details the symbolic reachability analysis and Section 4 outlines the implementation details. Section 5 provides empirical results from performing reachability analysis benchmarking examples in PROB alone and using the new integration of the two tools. The paper concludes in Section 6 with reflections and future work.

2 Preliminaries: B-Method and Transition Systems

In this section we provide an overview of the B-Method and the foundations used within LTSMIN.


A B machine consists of a collection of clauses and a collection of operations. The MACHINE clause declares the abstract machine and gives it its name. The VARIABLES clause declares the variables that are used to carry the state information within the machine. The INVARIANT clause gives the type of the variables, and more generally it also contains any other constraints on the allowable machine states. The INITIALISATION clause determines the initial state(s) of the machine. Operations in a machine are events that change the state of a machine and can have input parameters. Operations can be of the form SELECT P THEN S END where P is a guard and S is the action part of the operation. The predicate P must include the type of any input variables and also give conditions on when the operation can be performed. When the guard of an operation is true then the operation is enabled and can be performed. If the guard is the simple predicate *true* then the operation form is simplified to BEGIN S END. An operation can also be of the form PRE P THEN S END so that the predicate is a precondition and if the operation is invoked outside its precondition then this results in a divergence (we do not illustrate this in our running example). Finally, the action part of an operation is a *generalised substitution*, which can consist of one or more assignment statements (in parallel) to update the state or assign to the output variables of an operation. Conditional statements and nondeterministic choice statements are also permitted in the body of the operation. The example in Figure 2 illustrates the *MutexSimple* machine with three variables and five operations. Its initial state is deterministic and *wait* is set to MAXINT. For MAXINT=1 we get 4 states; the state space constructed by ProB can be found in Figure 3. From the initial state only the guards for **Enter** and **Leave** are true. Following an **Enter** operation the value of the **cs** variable is true which means that the guard of the **CS_Active** operation is true and the system can indicate that it is in the critical section by performing the **CS_Active** operation.

The example presented could also be considered as an **Event-B** example since it is a simple guarded system. We do not elaborate further on the notation of Event-B in this paper but note that the results in the subsequent sections are also applicable to Event-B.

```

1 MACHINE MutexSimple
2 VARIABLES cs, wait, finished
3 INVARIANT
4   cs: BOOL & wait: NATURAL & finished: NATURAL
5 INITIALISATION cs := FALSE || wait := MAXINT || finished := 0
6 OPERATIONS
7   Enter      = SELECT cs = FALSE & wait > 0 THEN
8             cs := TRUE || wait := wait - 1 END;
9   Exit       = SELECT cs = TRUE THEN
10            cs := FALSE || finished := finished + 1 END;
11   Leave     = BEGIN cs := FALSE END;
12   CS_Active = SELECT cs = TRUE THEN skip END;
13   Restart   = SELECT finished > 0 THEN
14            wait := wait + 1 || finished := finished - 1 END
15 END

```

Fig. 2: *MutexSimple* B-Method machine exampleFig. 3: *MutexSimple* statespace for MAXINT=1

As far as symbolic reachability analysis is concerned, a formal model is seen to denote a transition system. LTSMIN adopts the following definition:

Definition 1 (Transition System). A *Transition System (TS)* is a structure (S, \rightarrow, I) , where S is a set of states, $\rightarrow \subseteq S \times S$ is a transition relation and $I \subseteq S$ is a set of initial states. Furthermore, let \rightarrow^* be the reflexive and transitive closure of \rightarrow , then the set of reachable states is $R = \{s \in S \mid \exists s' \in I. s' \rightarrow^* s\}$.

A B-Method and Event-B model induces such a transition system: initial states are defined by the initialisation clause and the individual operations together define the transition relation \rightarrow . Figure 3 shows the transition system⁴ for the machine in Figure 2. As can be seen in Figure 3, the transition relation is annotated with operation names. For symbolic reachability analysis it is actually very important that we divide the transition relation into groups, leading to the concept of a partitioned transition system:

⁴One subtle issue is that LTSMIN actually only supports a single initial state; this is solved by introducing the artificial *root* state linked to the initial states proper. We ignore this technical issue in the paper.

Definition 2 (Partitioned Transition System). A *Partitioned Transition System (PTS)* is a structure $\mathcal{P} = (S^N, G, \rightarrow^M, I^N)$, where

- $S^N = S_1 \times \dots \times S_N$ is the set of states, which are vectors of N values,
- $G = (\rightarrow_1, \dots, \rightarrow_M)$ is a vector of M transition groups $\rightarrow_i \subseteq S^N \times S^N$ ($\forall 1 \leq i \leq M$)
- $\rightarrow^M = \bigcup_{i=1}^M \rightarrow_i$ is the overall transition relation induced by G , i.e., the union of the M transition groups, and
- $I^N \subseteq S^N$ is the set of initial states.

We write $s \rightarrow_i t$ when $(s, t) \in \rightarrow_i$ for $1 \leq i \leq M$, and $s \rightarrow^M t$ when $(s, t) \in \rightarrow^M$.

For example $I^N = \{(FALSE, MAXINT, 0)\}$ in the running example. Note that G in Definition 2 does not necessarily form a partition of \rightarrow^M , overlap is allowed between the individual groups.

3 Symbolic Reachability Analysis for B

Computing the set of reachable states (R) of a transition system can be done efficiently with symbolic algorithms if many transition groups \rightarrow_i touch only a few variables. This concept is known as event locality [8]. Many models of transition systems in the B-Method employ event locality. In the B-Method event locality occurs in operations, where only a few variables are read from, or written to. For example in Figure 2 operation `CS_Active` only reads from `cs` and `Leave` only writes to `cs`. This event locality benefits the symbolic reachability analysis, so that the algorithm is capable of coping with the well known state space explosion problem. Since the B-Method employs event locality we build on the foundations of earlier work on LTSMIN [6, 24] and extend it to PROB. To perform symbolic reachability analysis of the B-Method, PROB should provide LTSMIN with read matrices and write matrices. These matrices inform LTSMIN about the locality of events in the B-Method.

Read independence is an important concept, it allows one to reuse the successor states computed in one state s for all states s' which differ just by read-independent variables from s , and vice versa.

Definition 3 (Read independence). Two state vectors s, s' are equivalent except on index j , denoted by $s \approx_j s'$, iff $\forall k \neq j : s_k = s'_k$.

Transition group i is read-overwrite independent from state variable j , iff $\forall s, s', t \in S^N$ such that $s \approx_j s'$ and $s \rightarrow_i t$, we have that $s' \rightarrow_i t$.

Transition group i is read-copy independent from state variable j , iff $\forall s, s', t \in S^N$ such that $s \approx_j s'$ and $s \rightarrow_i t$, we have that $s' \rightarrow_i (t_1, \dots, t_{j-1}, s'_j, t_{j+1}, \dots, t_N)$.

A transition group is read independent iff it is either read-overwrite or read-copy independent.

If an event never reads but may write to a variable j it generally does not satisfy the above definition. For example, the operation `MayReset = IF cs = true THEN wait := 0 END` would neither be read-copy nor read-overwrite

independent (for state vectors with $cs = \mathbf{false}$ it satisfies the definition of the former and for $cs = \mathbf{true}$ the latter, but neither for all state vectors). LTSMIN can also deal with more liberal independence notions, but we have not yet implemented this in the present paper.

Definition 4 (Write independence). *Transition group i is write-independent from state variable j , if $\forall s, t \in S^N: (s_1, \dots, s_j, \dots, s_N) \rightarrow_i (t_1, \dots, t_j, \dots, t_N) \implies (s_j = t_j)$, i.e. state variable j is never modified by transition group i .*

We illustrate the above definitions below.

Definition 5 (Dependency Matrices). *For a PTS $\mathcal{P} = (S^N, G, \rightarrow^M, I^N)$, the write matrix is an $M \times N$ matrix $WM(\mathcal{P}) = WM_{M \times N}^{\mathcal{P}} \in \{0, 1\}^{M \times N}$, such that $(WM_{i,j} = 0) \implies$ transition group i is write independent from state variable j . Furthermore, the read matrix is an $M \times N$ matrix $RM(\mathcal{P}) = RM_{M \times N}^{\mathcal{P}} \in \{0, 1\}^{M \times N}$, such that $(RM_{i,j} = 0) \implies$ transition group i is read independent from state variable j .*

In this paper we will use sufficient syntactic conditions to ensure Definitions 3 and 4 and obtain the read and write matrix from Definition 5. Indeed, we compute for every operation syntactically which variables are read from and which variables are written to.

- If an operation does not write to a variable, its transition group is write independent according to Definition 4 and the corresponding entry in WM is 0.
- If an operation does not read a variable, its transition group is read independent according to Definition 3, unless it maybe written to (e.g., because the assignment is in the branch of an if-then-else). In this case, we will mark the variable as both write and read independent. Also, note that when the assignment within an operation is of the form $f(X) := E$ then the operation should have a read dependency on the function f (in addition to the write dependency).

For our example in Figure 2 the syntactic read-write information is as follows:

	cs	wait	finished		cs	wait	finished
Enter	1	1	0	Enter	1	1	0
Exit	1	0	1	Exit	1	0	1
Leave	0	0	0	Leave	1	0	0
CS_Active	1	0	0	CS_Active	0	0	0
Restart	0	1	1	Restart	0	1	1
(a) Read matrix (RM)				(b) Write matrix (WM)			

Fig. 4: Dependency matrices

From the matrices we can infer if a variable is read-copy or read-overwrite independent: a variable that is read independent and not written to (i.e., write independent) is read-copy independent, otherwise it is read-overwrite independent. We can thus infer that:

- the transition group of **Enter** is read-copy and write independent on **finished**.

- **Exit** is read-copy and write independent on **wait**.
- **Leave** is read-copy and write independent on **wait** and **finished** and read-overwrite independent on **cs**.
- **CS_Active** is read-copy and write independent on **wait** and **finished** and write independent on **cs** (but not read-independent on **cs**).
- **Leave** is read-copy and write independent on **cs**.

3.1 Exploration Algorithm

We now present the core of the symbolic reachability analysis algorithm of *LTSMIN*. Algorithm 1 computes the set of reachable states R (represented as a decision diagram) and it uses the independence information to minimise the number of next state computations that have to be carried out, i.e., re-using the next states $\{t \mid s \rightarrow_i t\}$ computed for a single state s for many other states s' according to Definitions 3 and 4. Algorithm 1 will, while it keeps finding new states, expand the partial transition relation with potential successor states, and apply the expanded relation to the set of new states.

Four key functions that make Algorithm 1 highly performant are the following.⁵ The (1) *read projection* $\pi_i^r = \pi_i^{RM}$ and (2) *write projection* $\pi_i^w = \pi_i^{WM}$ take as argument a state vector and produce a state vector restricted to the read and write dependent variables of group i , respectively. Furthermore these function are extended to apply to sets directly, e.g., given the examples in Figures 2 and 4, a read projection for **Leave** is $\pi_3^r(\{(FALSE, 0, 0), (FALSE, 0, 1), (FALSE, 1, 0)\}) = \{(FALSE)\}$. This is illustrated in Figure 6 and used at Line 2 in Algorithm 2. The read projection prevents *LTSMIN* from doing two unnecessary next state calls to **PROB**, since **Leave** is *read-copy* independent on **wait** and **finished**.

The function (3) NEXTSTATE_i takes a read projected state and projects (with π_i^w) all successor states of transition group i . The partial transition relation \hookrightarrow_i^P is learned on the fly, and NEXTSTATE_i is used to expand \hookrightarrow_i^P . An example next state call for **Enter** is $\text{NEXTSTATE}_1((FALSE, 1)) = \{(TRUE, 0)\}$.

Lastly, (4) **NEXT** takes a set of states, a partial transition relation, a row of the read and write matrix and outputs a set of successor states. For example, applying the partial relation of **Enter** to the initial state yields $\text{NEXT}(\{(FALSE, \underline{1}, 0)\}, \{((FALSE, 1), (TRUE, 0)), (1, 1, 0), (1, 1, 0)\}) = \{(TRUE, 0, 0)\}$. Note that in this example **Enter** is *read-copy* independent on **finished** and thus **NEXT** will copy its value from the initial state.

The usage of these four key functions is also illustrated in Figure 5. The figure shows that first the projection is done for **Enter**, then \hookrightarrow_i^P is expanded with a NEXTSTATE_i call, lastly relation \hookrightarrow_i^P is applied to the initial state, producing the first successor state.

Figure 6 shows for each operation the transition relation \hookrightarrow_i^P and the projected states on which they are computed. Definition 3 ensures that the projected state

⁵We refrain from giving their formal definitions; they can be found in [24].

Algorithm 1: REACHBREADTHFIRST

Input : $I^N \subseteq S^N, M \in \mathbb{N}, RM, WM$
Output: The set of reachable states \mathcal{R}

```

1  $\mathcal{R} \leftarrow I^N; \mathcal{L} \leftarrow \mathcal{R};$ 
2 for  $1 \leq i \leq M$  do  $\mathcal{R}_i^P \leftarrow \emptyset; \hookrightarrow_i^P \leftarrow \emptyset$  while
 $\mathcal{L} \neq \emptyset$  do
3   LEARNTRANS();  $\mathcal{N} \leftarrow \emptyset;$ 
4   for  $1 \leq i \leq M$  do
5      $\mathcal{N} \leftarrow \mathcal{N} \cup \text{NEXT}(\mathcal{L}, \hookrightarrow_i^P, RM_i, WM_i);$ 
6    $\mathcal{L} \leftarrow \mathcal{N} - \mathcal{R}; \mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{N};$ 
7 return  $\mathcal{R}$ 
```

Algorithm 2: LEARNTRANS

Description: Extends \hookrightarrow_i^P

```

1 for  $1 \leq i \leq M$  do
2    $\mathcal{L}^P \leftarrow \pi_i^r(\mathcal{L});$ 
3   for  $s^P \in \mathcal{L}^P - \mathcal{R}_i^P$  do
4      $\hookrightarrow_i^P \leftarrow \hookrightarrow_i^P \cup \{(s^P, d^P) \mid$ 
5        $d^P \in$ 
6        $\text{NEXTSTATE}_i(s^P)\};$ 
7    $\mathcal{R}_i^P \leftarrow \mathcal{R}_i^P \cup \mathcal{L}^P;$ 
```




Fig. 5: Illustrating one iteration of Algorithm 1 for MutexSimple and **Enter** operation

space shown in Figure 6 can be used to compute the effect of each of these operations for the *entire* state space (using **next**).

3.2 List Decision Diagrams

The symbolic reachability algorithm in Section 3.1 uses List Decision Diagrams (LDDs) to store the reachable states and transition relations. Similar to a Binary Decision Diagram, an LDD [6] represents a set of vectors. Due to the sharing of state vectors within an LDD, the memory usage can be very low, even for very large state spaces. Three example LDDs for the running example are given in Figure 7. The LDDs represent the set of reachable states \mathcal{R} in Algorithm 1 at each iteration of Line 2. In an LDD every path from the top left node to $\{\epsilon\}$ is a state, e.g., the initial state $(FALSE, 1, 0)$ in Figure 7b. A node in an LDD represents a unique set of (sub) vectors, e.g., $\{\epsilon\}$ represents the set of zero-length vectors and the right-most 0 of variable **wait** in Figure 7d encodes the set $\{(0, 0), (0, 1), (1, 0)\}$. Figure 7c shows that firing **Enter** will add $(TRUE, 0, 0)$ to R . In Figure 7d $(FALSE, 0, 0)$ and $(FALSE, 0, 1)$ are added to R , by firing **Leave** and **Exit** respectively. The benefit of using LDDs for state storage is due to the sharing of state vectors. For example, the subvector $(FALSE)$ of the states $\{(FALSE, 0, 0), (FALSE, 0, 1), (FALSE, 1, 0), (FALSE, 1, 1)\}$ in iteration 3 is encoded in the LDD with a single node. For bigger state spaces the sharing can be huge; resulting in a low memory footprint for the reachability algorithm.

3.3 Performance: NEXTSTATE function

There are two big differences of Algorithm 1 with classical explicit state model checking as used by PROB [20]. First, the state space is represented using an

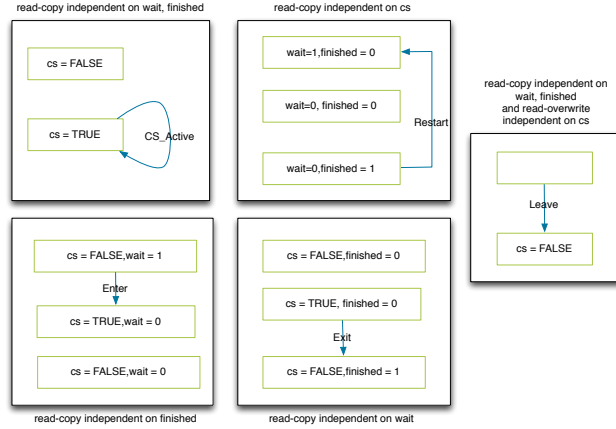


Fig. 6: MDDs of the reachable states

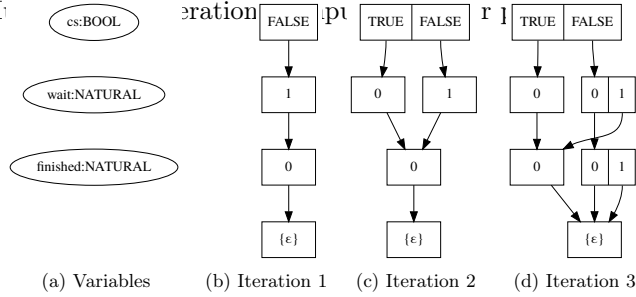


Fig. 7: LDDs of the reachable states

LDD datastructure, which enables sharing amongst states. Second, independence is used to apply the NEXTSTATE function not state by state, but for entire sets of states in one go. For each of the 4 states in Figure 3, the explicit model checking algorithm of PROB would check whether each of the 5 operations is enabled; resulting in 20 next-state calls. With LTSMIN’s symbolic reachability Algorithm 1, only 12 NEXTSTATE calls are made. This is shown in the following table, where + means enabled, - means disabled, and C means that LTSMIN has reused the results of a previous call to PROB.

State#	cs	wait	finished	Enter	Exit	Leave	CS_Active	Restart
1	FALSE	1	0	+	C	C	C	-
2	TRUE	0	0	-	+	+	+	-
3	FALSE	0	0	-	-	C	-	C
4	FALSE	0	1	C	-	C	C	+

If we initialise wait with MAXINT = 500, the state space has 251,002 states. The runtime with PROB is 70 seconds, with LTSMIN+PROB 48 seconds and LTSMIN performs only 6012 NEXTSTATE calls. The example does not have a lot of concurrency and uses only simple data structures (and thus the overhead of the LTSMIN’s PROB front-end is more of a factor compared to the runtime of

ProB for computing successor states); other examples will show greater speedups (see Section 5). But the purpose of this example is to illustrate the principles.

4 Technical Aspects and Implementation

We used a distributed approach to integrate ProB and LTSMIN. Both tools are stand-alone applications, so a direct integration, i.e., turning one of the tools into a shared library would require considerable effort. We therefore added extensions to both tools that convert the data formats and use sockets to communicate with each other. A high level view of the integration is shown in Figure 8. We use the $\text{\O}MQ$ [13] library for communication. $\text{\O}MQ$ is oriented around message queues and can be used as both, a networking library with very high throughput and as a concurrency framework. We have chosen $\text{\O}MQ$ because it worked very well in previous work [4]. Although we do not (yet) have to care about concurrency in this work, the reactor abstraction provided by $\text{\O}MQ$ was very handy in the ProB extension. It allows to implement a server that receives and processes messages without much effort. The communication is always initiated by LTSMIN; it sends a message and blocks until it receives the answer from ProB.

We usually run both tools on a single computer using interprocess (IPC) sockets, but it is only a matter of configuration to run the tools on different machines using TCP sockets. We currently only support Linux and Mac OS. The communication protocol is straightforward. Reachability analysis is initiated from LTSMIN by sending an initialisation packet. ProB answers with a message containing the relevant static information about a model, such as the dependency matrices that LTSMIN requires (see Section 3).

Each matrix is encoded as a 2-dimensional array, which is not optimal for a sparse matrix but is not an issue because we only send the matrices once. The packet also contains the list of variables, their types, the list of transition groups, and the initial state.




Fig. 8: High level design showing the integration

States are represented as a list of so called **chunks**. A chunk is one of the elements in the state vector according to Definition 2. In the case of B, each chunk is a value of one of the state variables. Because LTSMIN will not look inside the chunks, we simply use the binary representation of ProB's Prolog term that represents the value of a variable. This has the advantage, that ProB does not have to keep information about the state space. It can always recover a state from the chunks that are sent by LTSMIN. The transition groups correspond to B operations as explained in Figure 3. Like chunks the transition groups are only used as names in LTSMIN.

Once the initial setup is done, LTSMIN will start to ask ProB for successor states for **specific** transition groups. It will send a next-state message containing

a state and a transition group. The state, that LTSMIN sends is a list of chunks and PROB's LTSMIN extension can directly consume them and construct a Prolog term that internally represents a state. Using this constructed state and the transition group, the extension will then ask PROB for all successor states. The result is a list of Prolog terms, each representing a successor state. The extension transforms the list of states into a list of lists of chunks and sends them back to LTSMIN. This is repeated until LTSMIN has explored all necessary states and sends a termination signal.

The next-state messaging is similar to Figure 5, the projection is achieved by replacing all read independent variables by a dummy value.

5 Experiments

To demonstrate that the combination of PROB and LTSMIN improves the performance of the reachability analysis and deadlock detection compared with the standalone version of PROB, we use a wide range of B and Event-B models. Our benchmark suite contains puzzles (e.g., towers of Hanoi) as well as specifications of protocols (e.g., Needham-Schroeder), algorithms (e.g., Simpson's four slot algorithm) and industrial specifications (e.g., a choreography model by SAP, a cruise control system by Volvo and a fault tolerant automatic train protection system by Siemens).⁶

The experiments were run on Ubuntu 15.10 64-bit, with 8 GB RAM, 120 GB SSD and an Intel Sandybridge Mobile i5 2520M 2.50 GHz Dual core. The version of PROB used in this paper is 1.5.1-beta3, and LTSMIN tag LTSminProBiFM2016⁷

Figure 9 summarises a selection of the experiments that we ran. The last two models are Event-B models. In these experiments we used Breadth-First Search (BFS) and looked for deadlocks. A deadlock was found only for the Philosophers model (this is also why there are no next state call statistics for this model). The table also contains the number of next state calls for PROB reachability analysis on its own and when called from LTSMIN's symbolic reachability analysis algorithm (i.e., our new integration see Section 3.3) without deadlock checking. One can clearly see that we obtain a considerable reduction in wallclock time. The PROB time is the walltime of the PROB reachability analysis and initial state computation and does not include parsing and loading. The LTSMIN CPU time column shows how much time is spent in the LTSMIN side of the symbolic reachability analysis algorithm. The LTSMIN wall time shows the total walltime, and this also contains the time spent in the communication layer and waiting for the PROB process to compute the next states. To compare the benefit of our new algorithm we compute the speedup of the walltime in the last column by dividing the PROB walltime from column 5 with the LTSMIN walltime in column 7.

⁶More detailed descriptions can be found in Appendix A.

⁷The software and models can be found online at <https://github.com/utwente-fmt/ProB-LTSmin-1FM16>.

We can see that for some of the smaller models the overhead of setting up LTSMIN does not pay off. However, for all larger models, except for the Train1_Lukas_POR model considerable speedups were obtained.

Benchmark	Events	States	PROB Nxt St Calls	LTSMIN NxtSt Calls	PROB Wall (ms)	LTSMIN CPU (ms)	LTSMIN Wall (ms)	Speedup
CAN_BUS	21	132600	2784560	3534	122850	660	1590	77.264
ConcurrentCounters	4	110813	443249	113032	21820	2760	13820	1.579
Cruise_finite1	26	1361	35361	1667	2900	100	1020	2.843
file_system	8	698	5577	1198	1900	180	4660	0.41
MutexSimple	5	10	46	26	10	10	190	0.053
Philosophers	5				480	40	590	0.814
SiemensMiniPilot_Abrial0	9	181	1621	182	100	20	260	0.385
Simpson_Four_Slot	9	46658	419906	2089	17310	200	860	20.128
Train1_Lukas_POR	8	24637	197082	101441	33660	6480	50260	0.670
nota	11	80719	887899	588	287970	130	660	436.318
pkeyprot2	10	4412	44111	2004	22190	210	1710	12.977
Ref5_Switch_mch	38	29861	1134681	1281	160600	490	1260	127.460
obsw_M001	21	589279	12374779	23406	2051320	1620	12420	165.163

Fig. 9: B and Event-B Machines, with BFS and deadlock detection

A major result we achieved with non default settings for LTSMIN, is for elevator12.eventb. This model is not listed in Figure 9, because PROB runs out of memory on the hardware configuration used for this experiment. LTSMIN computed in 34 seconds, with 96,523 NEXTSTATE calls, that the model has 1,852,655,841 states. As reachability algorithm we chose chaining [30], and to compute a better variable order, we ran Sloan’s bandwidth reduction algorithm [33] on the dependency matrix.

As far as memory consumption is concerned; when performing reachability analysis on CAN_BUS, the PROB process consumes 370 MB real memory, while the LTSMIN process consumes 633 MB, with the default settings. With the default settings LTSMIN will allocate 2^{22} elements (≈ 100 MB) for the node table and 2^{24} elements (≈ 500 MB) for the operations cache. If we choose a smaller node table and operations cache for the LDD package (both 2^{18} elements), LTSMIN consumes only 22 MB. The default settings for LTSMIN are geared towards larger symbolic state spaces than that of CAN_BUS. The default node table and cache are too big for CAN_BUS, and thus not completely filled during reachability.

We have also run our new symbolic reachability analysis on Z and TLA⁺ models. For example, we successfully validated the video rental Z model from [9]. For 2 persons and 2 titles and maximum stock level of 4, LTSMIN generates the 23009 states in 1.75 seconds compared to 52.4 seconds with PROB alone. The model contained useless constants; after removing them PROB runs in 1.6 seconds; the runtime of LTSMIN stays unchanged. We were unable to use the output of z2sal [9] using SAL [26] and its symbolic model checker for comparison.

In summary, Figure 9 shows that for several non-trivial B and Event-B models, considerable improvements can be obtained using the symbolic reachability analysis technique described in this paper.

Alternate Approaches Other techniques for improving model checking for B-Method and Event-B models have been developed and evaluated in the recent years. We have run a further set of experiments using a selection of those methods; the complete results can be found in Appendix B. For technical reasons, the experiments were run on different hardware than above, a MacBook Air with 2.2 GHz i7 processor and 8GB of RAM. We summarise the findings here and compare the results with our new symbolic model checking algorithm.

Benchmark	PROB POR		PROB Hash		TLC		PROB no opt. ms
	ms	Speedup	ms	Speedup	sec	Speedup	
CAN_BUS	138720	0.80	98390	1.12	3	37	110400
ConcurrentCounters	50	345.8	18400	1.06	1	17	17290
file_system	2380	0.37	210	4.24	29	0.03	890
Simpson_Four_Slot	20860	0.70	9550	1.52	1	15	14530
Train1_Lukas_POR	34030	0.75	28930	0.88	4	6	25740
nota	490	509.22	14780	16.88	10	25	249520
Ref5_Switch_mch	215160	0.59	124500	1.01	6	21	126170
obsw_M001	2150520	0.80	76190	22.53	55	31	1716770

The authors in [11] presented a translation from the B-Method to TLA⁺, with the goal of using the **TLC** model checker [36] as backend. TLC has no constraints solving capabilities, and as such that it can only deal with lower level models. On the other hand, its execution can be considerably faster than PROB, and its explicit state model checking engine (which stores fingerprints) is very efficient. On the downside, there is a small probability that fingerprint collisions can occur. The experiments show that TLC does not deal well with benchmark programs which require constraint solving (graph isomorphism, JobsPuzzle, ...), running up to three orders of magnitude slower than PROB or LTSMIN with PROB. However, it does deal very well with lower level models, e.g., it is faster than LTSMIN for ConcurrentCounters. For many benchmark models, even those not requiring constraint solving, our symbolic reachability analysis is faster. For example, for the nota example, TLC runs in about 10 seconds—faster than PROB without any optimisation—but slower than LTSMIN by less than a second.

Symmetry reduction [21] can be very useful; but exponential improvements usually occur only on academic examples. Here we have experimented with the hash marker symmetry reduction, which is PROB’s fastest symmetry method, but is generally not guaranteed to explore all states. The method gives the best results for certain models (e.g., file_system). But for several of the larger, industrial examples shown above, its benefit is not of the same scale as LTSMIN. In future, we will investigate combining PROB’s symmetry reduction with the new LTSMIN algorithm.

We have also experimented with **partial order reduction**. [10] uses a semantic preprocessing phase to determine independence (different from our purely syntactic determination; see Section 3). As such, it can induce a slow down for some examples where this does not pay off (e.g., file_system). PROB’s partial order reduction obtains the best times for certain models with a large degree of concurrency (ConcurrentCounters, SiemensMiniPilot_Abrial, and nota). However, once we start doing invariant checking, [10] does not scale nearly as well (e.g., it

takes 134 seconds instead of 0.5 seconds for LTSMIN checking the nota model). But even without invariant checking, there are plenty examples where the symbolic reachability analysis approach is better (e.g., `Cruise_finite1`, `Philosophers`, `Simpson_Four_Slot` and almost two orders of magnitude for `CAN_BUS`). In summary:

- TLC is good for models not requiring constraint solving. It is a very efficient, explicit state model checker. However, models often have to be rewritten (such as `CAN_BUS`), and there is a small chance of having fingerprint collisions.
- Symmetry reduction excels when models make use of deferred sets. However, the hash marker method is not guaranteed to explore all states.
- Partial order reduction is very good for models with a large degree of concurrency. However, it can cause slow downs and is less suited for invariant checking.
- The new symbolic reachability analysis algorithm deals well with concurrency and is by far the fastest method for certain larger, industrial models, such as `CAN_BUS`, `obsw_M001`, `elevator12`, the ABZ landing gear model or Abrial’s mechanical press. LTSMIN is currently the only tool set that uses a symbolic representation of the state space that is connected to PROB.

6 More Related Work, Future Work and Conclusion

We have already evaluated the use of TLC [36] for model checking B. Another explicit state model checker for B has been presented in [22], which uses lazy enumeration. Symbolic model checking [7] has been used for railway applications in [35]. The best known symbolic model checker is probably SMV [23], which uses a low-level input language. Some comparisons between using SMV and PROB have been conducted in [14], where models were translated by hand. For abstract state machines there is the AsmetaSMV [3] tool, which automatically translates ASM specifications to SMV. It is our impression that the translation often leads to a considerable blowup of the model, encoded in SMV’s low-level language, also affecting performance. We did one experiment on a Tic-Tac-Toe model provided for AsmetaSMV: NuSMV 2.6 took over 13 seconds to find a configuration where the cross player wins; PROB (without LTSMIN) took 0.2 seconds model checking time for the same property on a similar B model. Another experiment involved puzzle 3 of the RushHour game: PROB solves this in 5 seconds, while NuSMV still had not found a solution after 120 minutes.

Other symbolic model checkers that perform comparable well to LTSMIN include MARCIE [12] and PETRIDOTNET [25].

The paper provides a stable architectural link between PROB and LTSMIN that can be extended. First, we plan to provide LTSMIN with more fine-grained information about the models, both statically and dynamically. Dynamically, PROB will transmit to LTSMIN which variables have actually been written by an operation, enabling a more extensive independence notion to be used. Statically, PROB will transmit the individual guards of operations and provide variable read matrices for the guards. We will also transmit the individual invariants in the

same manner, to enable analysis of the invariants. (It is actually already possible to check invariants using the present integration, simply by encoding invariants as operations. We have done so with success for some of the examples, e.g., the nota from Section 5.) When PROB transmits individual guards, we also hope to use the guard-based partial order optimisations of LTSMIN [17] and enable LTL model checking with LTSMIN.

These future directions will strengthen the capability of the verification tools and hence further encourage the application of formal methods within industry as identified in [5], for example to support complex railway systems verification in CSP||B. This will require both more fine-grained static and dynamic information.

In summary, we have presented a new scalable, symbolic analysis algorithm for the B-Method and Event-B, along with a platform to integrate other model checking improvements via LTSMIN in the future.

References

1. Abrial, J.R.: Modeling in Event-B: System and Software Engineering. Cambridge University Press (2010)
2. Abrial, J.: The B-book - assigning programs to meanings. Cambridge University Press (2005)
3. Arcaini, P., Gargantini, A., Riccobene, E.: Asmetasmv: A way to link high-level ASM models to low-level nusmv specifications. In: Proceedings ABZ'10. pp. 61–74 (2010)
4. Bendisposto, J.: Directed and Distributed Model Checking of B Specifications. Ph.D. thesis, University of Düsseldorf (2015), <http://docserv.uni-duesseldorf.de/servlets/DocumentServlet?id=34472>
5. Bicarregui, J., Fitzgerald, J.S., Larsen, P.G., Woodcock, J.C.P.: Industrial practice in formal methods: A review. In: FM 2009: Formal Methods, Second World Congress, Eindhoven, The Netherlands, November 2-6, 2009. Proceedings. pp. 810–813 (2009)
6. Blom, S., van de Pol, J.: Symbolic Reachability for Process Algebras with Recursive Data Types. In: Theoretical Aspects of Computing - ICTAC 2008, 5th International Colloquium, Istanbul, Turkey, September 1-3, 2008. Proceedings. pp. 81–95 (2008)
7. Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.J.: Symbolic model checking: 10^{20} states and beyond. Information and Computation 98(2), 142–170 (1992)
8. Ciardo, G., Marmorstein, R.M., Siminiceanu, R.: The saturation algorithm for symbolic state-space exploration. STTT 8(1), 4–25 (2006)
9. Derrick, J., North, S., Simons, A.J.H.: Z2SAL - building a model checker for Z. In: Börger, E., Butler, M., Bowen, J.P., Boca, P. (eds.) Proceedings ABZ 2008. pp. 280–293. LNCS 5238 (2008)
10. Dobrikov, I., Leuschel, M.: Optimising the ProB model checker for B using partial order reduction. In: SEFM 2014. pp. 220–234. LNCS 8702, Grenoble (2014)
11. Hansen, D., Leuschel, M.: Translating B to TLA + for validation with TLC. In: Proceedings ABZ'14. pp. 40–55. LNCS 8477 (2014)
12. Heiner, M., Rohr, C., Schwarick, M.: MARCIE - Model checking And Reachability analysis done effiCIEntly. In: Colom, J., Desel, J. (eds.) Proc. PETRI NETS 2013. LNCS, vol. 7927, pp. 389–399 (June 2013)
13. Hintjens, P.: ZeroMQ: Messaging for Many Applications. O'Reilly Media, Inc. (2013)

14. Hörne, T., van der Poll, J.A.: Planning as model checking: the performance of ProB vs NuSMV. In: Botha, R., Cilliers, C. (eds.) SAICSIT Conf. ACM International Conference Proceeding Series, vol. 338, pp. 114–123. ACM (2008)
15. James, P., Moller, F., Nguyen, H.N., Roggenbach, M., Schneider, S., Treharne, H., Trumble, M., Williams, D.M.: Verification of scheme plans using CSP || B. In: SEFM 2013 Collocated Workshops. pp. 189–204 (2013)
16. Kant, G., Laarman, A., Meijer, J., van de Pol, J., Blom, S., van Dijk, T.: LTSmin: High-performance language-independent model checking. In: TACAS 2015, London, UK, April 11-18, 2015. Proceedings. pp. 692–707 (2015)
17. Laarman, A., Pater, E., van de Pol, J.C., Hansen, H.: Guard-based partial-order reduction. *STTT* (2014)
18. Ladenberger, L., Hansen, D., Wiegard, H., Bendisposto, J., Leuschel, M.: Validation of the ABZ landing gear system using prob. *International Journal on Software Tools for Technology Transfer* pp. 1–17 (2015)
19. Legeard, B., Peureux, F., Utting, M.: Automated boundary testing from Z and B. In: Eriksson, L.H., Lindsay, P. (eds.) Proceedings FME’02. pp. 21–40. LNCS 2391, Springer-Verlag (2002)
20. Leuschel, M., Butler, M.J.: ProB: an automated analysis toolset for the B method. *STTT* 10(2), 185–203 (2008)
21. Leuschel, M., Massart, T.: Efficient approximate verification of B via symmetry markers. *Annals of Mathematics and Artificial Intelligence* 59(1), 81–106 (2010)
22. Matos, P.J., Fischer, B., Silva, J.P.M.: A lazy unbounded model checker for event-b. In: Breitman, K., Cavalcanti, A. (eds.) ICFEM. pp. 485–503. LNCS 5885 (2009)
23. McMillan, K.L.: Symbolic Model Checking. Ph.D. thesis, Boston (1993)
24. Meijer, J., Kant, G., Blom, S., van de Pol, J.: Read, Write and Copy Dependencies for Symbolic Model Checking. In: HVC 2014, Haifa, Israel, November 18-20, 2014. Proceedings. pp. 204–219 (2014)
25. Molnár, V., Darvas, D., Vörös, A., Bartha, T.: Saturation-Based Incremental LTL Model Checking with Inductive Proofs. In: TACAS Proceedings. pp. 643–657 (2015)
26. de Moura, L., Owre, S., Shankar, N.: The SAL language manual. Tech. rep., SRI International (2003), technical Report SRI-CSL-01-02 (Rev.2)
27. Needham, R., Schroeder, M.: Using encryption for authentication in large networks of computers 21(12), 993–999 (Dezember 1978)
28. Oliver, I.: Experiences in using B and UML in industrial development. In: Julliand, J., Kouchnarenko, O. (eds.) B’2007. pp. 248–251. LNCS 4355, Springer-Verlag (2007)
29. Plagge, D., Leuschel, M.: Seven at a stroke: LTL model checking for high-level specifications in B, Z, CSP, and more. *STTT* 11, 9–21 (2010)
30. Roig, O., Cortadella, J., Pastor, E.: Verification of Asynchronous Circuits by BDD-based Model Checking of Petri Nets. In: Proceedings Application and Theory of Petri Nets. pp. 374–391 (1995)
31. Scheuermann, B., Lochert, C., Mauve, M.: Implicit hop-by-hop congestion control in wireless multihop networks. *Ad Hoc Networks* (2007), doi: 10.1016/j.adhoc.2007.01.001
32. Schneider, S., Treharne, H.: CSP theorems for communicating B machines. *Formal Asp. Comput.* 17(4), 390–422 (2005), <http://dx.doi.org/10.1007/s00165-005-0076-7>
33. Sloan, S.W.: A FORTRAN program for profile and wavefront reduction. *International Journal for Numerical Methods in Engineering* 28(11), 2651–2679 (1989)
34. Whitefield, J.: Linking ProB and LTSmin (2015), Final Year Dissertation, University of Surrey

35. Winter, K.: Optimising ordering strategies for symbolic model checking of railway interlockings. In: Margaria, T., Steffen, B. (eds.) ISoLA (2). LNCS, vol. 7610, pp. 246–260. Springer (2012)
36. Yu, Y., Manolios, P., Lamport, L.: Model checking TLA⁺ specifications. In: Pierre, L., Kropf, T. (eds.) Proceedings CHARME'99. pp. 54–66. LNCS 1703 (1999)

A Models for experiments (for referees)

These model descriptions will be uploaded to a website, also containing the publicly available models.

The classical B models shown in the tables in Section 5 are:

- **CAN_BUS**
A model of a controller area network (CAN) bus developed by Colley.
- **ConcurrentCounters**
This is a very simple model with three concurrent counters that can be incremented independently.
- **Cruise_finite1**
Volvo Vehicle Function. The B specification machine has 15 variables, 550 lines of B specification, and 26 operations. The invariant consists of 40 conjuncts. This B specification was developed by Volvo as part of the European Commission IST Project PUSSEE (IST-2000-30103).
- **file_system** is a simple model of a file system with users, groups and access control.
- **MutexSimple**
This is the model shown in Fig. 2 of the paper.
- **Philosophers**
This is a B model of the well-known dining philosophers problem.
- **SiemensMiniPilot_Abrial0**
This Siemens Mini Pilot was developed within the Deploy Project. It is a specification of a fault-tolerant automatic train protection system that ensures that only one train is allowed on a part of a track at a time. The model contains a single refinement level and rather complex invariants. The model was translated to classical B from Event-B.
- **Simpson_Four_Slot**
A model of Simpson's four slot algorithm. This B model only represents the individual steps of the algorithm. It is intended to be used in conjunction with a CSP model to describe the sequencing of the steps. Here, the B model on its own is model checked (thus leading to invariant violations).
- **Train1_Lukas_POR**
This is the first level of refinement of the railway interlocking model in Chapter 17 of [1]. It uses a simplified topology and routes are released immediately when all blocks are free (to reduce the state space).
- **nota**
A model developed by Nokia within the RODIN Project⁸ for the validation and verification of Nokia's NoTA hardware platform; see [28].
- **pkeyprot2 (Needham-Schroeder)**
The Needham-Schroeder public key protocol is an authentication protocol for

⁸<http://rodin.cs.ncl.ac.uk/>

creating a secure connection over a public network [27]. The model consists of a network with the two normal users called Alice and Bob, an attacker named Eve and the keyserver. The first version of this protocol, developed in 1978, contains an error which was found in 1995 by Lowe. This model is a slightly simplified version (reducing the messages sent by Eve).

The Event-B models shown in the tables in Section 5 are:

- **Ref5_Switch_mch** This is the fifth level of refinement of a solution [18] to the ABZ'14 landing gear challenge.
- **obsw1_M001**
The Space Systems Finland example is a model of a subsystem used for the ESA BepiColombo mission. The model is a specification of parts of the BepiColombo On-Board software, that contains a core software and two subsystems used for telecommand and telemetry of the scientific experiments, the Solar Intensity X-ray and particle Spectrometer (SIXS) and the Mercury Imaging X-ray Spectrometer (MIXS). The model was a mini pilot of the Deploy project.
- **elevator12**
This is the twelfth refinement of an elevator model by ETH Zürich.

Here are additional experiments that were run, but whose results are not shown in the main paper (due to page limit restrictions).

- **BlocksWorldGeneric6**
A model of blocksworld, with six blocks; the goal being to put all blocks in the right-order on top of each other.
- **Echo**
The Echo algorithm is designed to find the shortest paths in a network topology.
- **Hanoi6**
The well-known towers of Hanoi puzzle with 6 discs.
- **scheduler_bztt**
The process scheduler from [19].
- **RushHour**
The Rush Hour puzzle.⁹ This is the hardest puzzle (number 40 in the regular version of the game). The shortest solution needs 83 moves; here we have explored the full state space of 4782 states and 29890 transitions.
- **Cansell_Contention**
A Firewire-Leader election protocol by Dominique Cansell.
- **CXCC0**
CXCC (Cooperative Crosslayer Congestion Control) [31] is a cross-layer approach to prevent congestion in wireless networks. The invariants used in the model are rather complex.
- **press_7b_mch**
This is the seventh level of refinement of Abrial's model of a mechanical press.

⁹See [http://en.wikipedia.org/wiki/Rush_Hour_\(board_game\)](http://en.wikipedia.org/wiki/Rush_Hour_(board_game)).

B More Experiments: Partial order reduction, Symmetry and TLC (for referees)

We have run a further set of experiments using a selection of alternate methods for improving model checking for B and Event-B models. These experiments were run on a MacBook Air with 2.2 GHz i7 processor. For the final version of the paper, these experiments will be re-run on the same hardware as in Sect. 5, for all relevant PROB and LTSMIN settings. The experiments, models and setup will be put onto a website for open access.

- [11] presented a translation from classical B to TLA+, with the goal of using the TLC . We used this to run TLC with one worker, no invariant, no deadlock and no assertion and no LTL checking. The table in Sect. 5 of the paper contains just the TLC model checking time (as measured by TLC to a one second accuracy). However, the table below includes the full runtime; including JVM startup and parsing. The factor removes 1 second from runtime; in the final version of the paper we will obtain more precise model checking times for TLC. It can also not be directly applied to Event-B or Z models (requiring a translation to classical B first); hence the numbers for the Event-B models are missing.
- Hash marker symmetry reduction [21]. It is the fastest symmetry method, but is generally not guaranteed to explore all states The time is the walltime of the model checking and initial state computation and does not include parsing and loading.
- partial order reduction method of [10]. It uses a semantic preprocessing phase to determine independence. The time is the walltime of the model checking and initial state computation and does not include parsing and loading.

An analysis of these results can be found in Section 5. First, we present the models also shown in Figure 9 in Section 5:

Benchmark	POR		Hash		TLC		No opt.
	ms	Factor	ms	Factor	ms	Factor	ms
CAN_BUS	138720	1.26	98390	0.89	5340	0.04	110400
ConcurrentCounters	50	0.00	18400	1.06	2130	0.07	17290
Cruise_finite1	5060	1.92	2330	0.89	2950	0.74	2630
MutexSimple	30	1.50	30	1.50	1660	33.00	20
Philosophers	4180	52.25	100	1.25	1690	8.63	80
SiemensMiniPilot_Abrial_0	70	0.70	100	1.00	1440	4.40	100
Simpson_Four_Slot	20860	1.44	9550	0.66	2670	0.11	14530
Train1_Lukas_POR	34030	1.32	28930	1.12	6650	0.22	25740
nota	490	0.00	14780	0.06	11780	0.04	249520
pkeyprot2	24140	1.29	17990	0.96	1190	0.01	18770
Ref5_Switch_mch	215160	1.71	124500	0.99	-	-	126170
obs_w_M001	2150520	1.25	76190	0.04	-	-	1716770

These are some additional experiments that we performed; TLC could not run on the RushHour model due to a bug in TLC.

Benchmark	POR		Hash		TLC		No opt. ms
	ms	Factor	ms	Factor	ms	Factor	
BlocksWorldGeneric6	2980	0.94	950	0.30	2240	0.39	3160
CAN_BUS_normalized	105460	1.43	77090	1.05	4590	0.05	73530
CSM	700	1.37	50	0.10	1790	1.55	510
EchoAlg	1000	3.57	280	1.00	2730	6.18	280
GardnerSwitchingPuzzle_v2	1390	1.78	780	1.00	4510	4.50	780
GraphIsoMedium	50	0.83	110	1.83	40360	656.00	60
Hanoi6	440	0.98	440	0.98	1540	1.20	450
JobsPuzzle	100	1.11	80	0.89	32410	349.00	90
NQueens	40	1.00	30	0.75	20550	488.75	40
PhilRing	1500	1.95	120	0.16	1820	1.06	770
Reading	120	1.00	40	0.33	1470	3.92	120
brueckenproblem	530	1.89	300	1.07	2150	4.11	280
fahrzeugverwaltung2	2880	1.18	90	0.04	2840	0.75	2440
file_system	2380	2.67	210	0.24	28600	31.01	890
scheduler_bztt	260	3.25	60	0.75	1340	4.25	80
RushHour_v2_TLC	18670	1.25	16520	1.11	err	-	14920
cont0	330	0.65	320	0.63	1290	0.57	510
Model_4_NoDeadlock_v6	850	0.50	590	0.35	-	-	1710
VM_4_mch	1050	1.62	860	1.32	-	-	650
cxcc0	410	5.86	70	1.00	-	-	70
press_7b_mch	9960	1.23	2500	0.31	-	-	8090
scheduler	50	1.25	30	0.75	-	-	40