

# Proof Assisted Symbolic Model Checking for B and Event-B

Sebastian Krings<sup>(✉)</sup> and Michael Leuschel

Institut für Informatik, Universität Düsseldorf,  
Universitätsstr. 1, 40225 Düsseldorf, Germany  
{krings,leuschel}@cs.uni-duesseldorf.de

**Abstract.** We have implemented various symbolic model checking algorithms, like BMC, k-Induction and IC3 for B and Event-B. The high-level nature of B and Event-B accounts for complicated constraints arising in these symbolic analysis techniques. In this paper we suggest using static information stemming from proof obligations to simplify occurring constraints. We show how to include proof information in the aforementioned algorithms. Using different benchmarks we compare explicit state to symbolic model checking as well as techniques with and without proof assistance. In particular for models with large branching factor, e.g., due to complicated data values being manipulated, the symbolic techniques fare much better than explicit state model checking. The inclusion of proof information results in further clear performance improvements.

**Keywords:** B-Method · Event-B · Proof · Symbolic model checking

## 1 Introduction and Motivation

Model checking is one of the key techniques used in formal software development. Two variants are currently in use: explicit state model checking and symbolic model checking. In explicit state model checking, every state is computed, the invariant is verified and discovered successor states are queued to be analyzed themselves. Symbolic model checking on the other hand tries to represent the state space and possible paths through it by predicates representing multiple states at once. Instead of stepwise exploration of the state space graph, the model checking problem is encoded as a formula and given to a constraint solver.<sup>1</sup>

So far, existing model checkers for B and Event-B like PROB [20,21], Eboc [23], pyB [30] or TLC [31] (via [17]) rely on explicit state model checking. PROB features some symbolic techniques for error detection [14] and test-case generation [27], but not full-blown symbolic model checking. This is mostly due

---

S. Krings and M. Leuschel—Part of this research has been initially sponsored by the EU funded FP7 project 287563 (ADVANCE).

<sup>1</sup> BDD-style model checking [10] is also called symbolic model checking. In recent work PROB has been integrated with LTSMIn for such kind of model checking.

to the high-level nature of B and Event-B. Both the usage of higher-order constructs and the underlying non-determinism accounts for complicated constraints during symbolic model checking. Some complexity can be coped with by relying on SMT solvers [12] or SAT solvers [25], but this is not always the case.

In this paper we have implemented various symbolic model checking algorithms for B and then study various ways to use proof information to optimize them. The proof information is used to strengthen constraints and reduce the counterexample search space. For Event-B, the information stems from discharged proof obligations exported from Rodin [1]. For classical B, no automatic proof information is available at the moment within PROB.<sup>2</sup> However, we can recompute the information using PROB’s proof capabilities as outlined in [18]. Essentially, for a B operation with before-after-predicate  $BA$  we search for a solution to

$$\text{invariant}(x) \wedge BA(x, x') \wedge \neg \text{conjunct\_of\_invariant}(x').$$

If PROB reports a contradiction, we know that the operation can not lead to a violation of the particular conjunct. In addition to the technique in [18], we developed a bridge to the Atelier B provers ml and pp.

All techniques used in this paper have been implemented both for classical B and Event-B. Both languages will be used in our empirical evaluation. For the sake of brevity we will only talk about Event-B events in the following sections instead of distinguishing events and operations.

We will introduce the model checking algorithms BMC, k-Induction and IC3 in Sects. 2.1, 2.2, and 2.3. For each of them we will show how to include proof information into the occurring constraints. Following, in Sect. 3, we will empirically compare symbolic model checking to explicit state model checking and model checking with and without proof assistance. Discussion and conclusions will be presented in Sect. 4.

## 2 Proof Assisted Symbolic Model Checking

When using the B method to develop a software or system, one often alternates between different phases. Among those are writing and adapting the specification, manual and automated proof efforts as well as model checking. Yet, the different steps are only loosely coupled when it comes to tool support.

In [4] the authors have shown how to augment explicit state model checking with proof information. In the following, we introduce three symbolic model checking techniques for B and incorporate proof information in a similar fashion.

We will use the running example in Fig. 1 to illustrate various concepts in our paper. First, let us introduce the notation we will be using. By  $x$  we will denote a vector of state variables.  $x'$  denotes the state variables in the successor state. A predicate  $p$  over the state variables  $x$  is denoted by  $p(x)$ . The same predicate over the successor state is written as  $p(x')$ . By *Events* we denote the set of events of a B machine. By *Inv* we denote its invariant.

<sup>2</sup> In theory, one could export proof information from Atelier B.

```

MACHINE Counter
CONSTANTS m
PROPERTIES m : {127,255}
VARIABLES c
INVARIANT c>=0 & c<=m
INITIALISATION c:=0
OPERATIONS
    incby(i) = PRE i:1..64 THEN c := c+i END
END
    
```

**Fig. 1.** A simple, erroneous B machine

**Definition 1.** For an event  $evt \in \text{Events}$  let  $BA_{evt}(x, x')$  denote the before-after-predicate connecting state variables in  $x$  to their successors in  $x'$ .

**Definition 2.** For a predicate  $p = \bigwedge_{i \in I} p_i$  and event  $evt$  let  $proven_{evt,p}$  denote a set of conjuncts  $p_i$  that are proven to hold after the execution of  $evt$  on a  $p$ -state, i.e., we have  $proven_{evt,p} = \bigwedge_{i \in J} p_i$  for some  $J \subseteq I$  such that

$$\forall x, x'. p(x) \wedge BA_{evt}(x, x') \Rightarrow proven_{evt,p}(x').$$

Let  $unproven_{evt,p} = \bigwedge_{i \in I \setminus J} p_i$  denote the complement of  $proven_{evt,p}$ , i. e., all the conjuncts of  $p$  that are not in  $proven_{evt,p}$ . We also define  $proven_{evt} = proven_{evt,Inv}$  and  $unproven_{evt} = unproven_{evt,Inv}$ .

For the example in Fig. 1, we have  $BA_{incby}(c, c') = \exists i \in 1..64 \wedge c' = c + i$ . Furthermore,  $proven_{incby}(c) = c \geq 0$ ; the invariant  $c \geq 0$  is preserved by  $incby$ . This implies  $unproven_{incby}(c) = c \leq m$ .

In our current implementation, we have that  $proven_{evt,p} = \bigwedge_{j \in J} p_j$  with  $J \subseteq I$ . This however is not a strict limitation. One could add other predicates discovered to be implied to  $proven_{evt,p}$  so as to further strengthen the predicates given below.

**Lemma 1.** A valid solution for Definition 2 is always  $proven_{evt,p} = true$ , meaning that nothing is proven for the event  $evt$ . At the other extreme, if all conjuncts of  $p$  are proven to hold after the execution of  $evt$  then  $proven_{evt,p} = p$  and  $unproven_{evt,p} = true$ .

**Definition 3.** By  $T$  we refer to a monolithic transition predicate, i. e., the disjunction of all before-after-predicates:  $T(x, x') = \bigvee_{e \in \text{Events}} BA_e(x, x')$ . By  $I$  we denote the after predicate of the initialization; including the properties about the constants.

For Fig. 1 we have  $I(c) = m \in \{127, 255\} \wedge c = 0$ .

In the following sections, we show how proof information can be embedded in the queries of bounded model checking (Sect. 2.1), k-Induction based model checking (Sect. 2.2) and IC3 (Sect. 2.3). An empirical evaluation of the algorithms and the influence of using proof information will be performed in Sect. 3.

## 2.1 BMC — Bounded Model Checking

BMC [5] has been suggested by Armin Biere et al. in 1999 [6]. One of the main goals is to avoid the blowup and resulting slow down of BDDs-based model checking algorithms. This is achieved by replacing the BDDs by a SAT solver.

The basic idea is as follows: For an initial state relation  $I$ , a transition relation  $T$  and a property  $p$  and a bound  $k$  starting with  $k = 0$ , a sequence of propositional formulas is generated. Each of the formulas is satisfiable if and only if there exists a counterexample to the property with length  $\leq k$ . This can be expressed as:

$$BMC(p, k) = I(s_0) \wedge \bigwedge_{i=0}^{k-1} T(s_i, s_{i+1}) \wedge \bigvee_{i=0}^k \neg p(s_i)$$

In order to include proof information we have to rewrite the predicate. First of all, if we increase  $k$  step-by-step as done in PROB's implementation of BMC, it is sufficient to check only the last state for a violation of  $p$ :

$$BMC(p, k) = I(s_0) \wedge \bigwedge_{i=0}^{k-1} T(s_i, s_{i+1}) \wedge \neg p(s_k)$$


For the example machine in Fig. 1, we have:

$$BMC(Inv, 0) = m \in \{127, 255\} \wedge c_0 = 0 \wedge \neg(c_0 \geq 0 \wedge c_0 \leq m) \quad (1)$$


$$BMC(Inv, 1) = m \in \{127, 255\} \wedge c_0 = 0 \wedge \exists i.(i \in 1..64 \wedge c_1 = c_0 + i) \\ \wedge \neg(c_1 \geq 0 \wedge c_1 \leq m) \quad (2)$$

$$BMC(Inv, 2) = m \in \{127, 255\} \wedge c_0 = 0 \wedge \exists i.(i \in 1..64 \wedge c_1 = c_0 + i) \\ \wedge \exists i.(i \in 1..64 \wedge c_2 = c_1 + i) \wedge \neg(c_2 \geq 0 \wedge c_2 \leq m) \quad (3)$$

PROB's constraint solver finds no solution for Eqs. (1) and (2), but does so for Eq. (3):  $m = 127, c_0 = 0, c_1 = 64, c_2 = 128$ . One can see that the constraint solver has instantiated the parameter of the event in such a way as to violate the invariant. PROB's classical model checker on the other hand "blindly" enumerates all 64 possible successor states. Using breadth-first search, the counterexample is found after having generated 325 states and 12420 transitions; taking  $\sim 1.5s$  whereas BMC finds the counterexample with  $k = 2$ , i.e., after three calls to the constraint solver and  $\sim 1s$ . A depth-first search may generate a long counterexample of up to 127 steps, depending in which order the successors are processed. PROB in this case actually processes the successors with the larger  $i$  values first; leading to a counterexample of length 4 after generating 324 states and 323 transitions. The state space is shown in Fig. 2, the corresponding counterexample is shown in Fig. 3. The larger the branching-factor, the better BMC becomes as compared to explicit state model checking. When the number



**Fig. 2.** State space of explicit state model checking



**Fig. 3.** Counterexample found by BMC

of possible parameter values becomes unbounded, e.g., supposing the `incby` event had no upper bound on  $i$ , BMC is often the only practical solution.<sup>3</sup>

Next, we extend the transition relation to either assert a property after every step or assert its negation:

**Definition 4.** For a predicate  $p$  we define  $T_p$  and  $T_p^-$  by

$$T_p(x, x') = \bigvee_{e \in Events} (BA_e(x, x') \wedge p(x'))$$

$$T_p^-(x, x') = \bigvee_{e \in Events} (BA_e(x, x') \wedge proven_{evt,p}(x') \wedge \neg unproven_{evt,p}(x'))$$

For  $k \geq 1$ , the proven conjuncts of  $p$  can be used to strengthen the constraint:

$$BMC(p, k) = I(s_0) \wedge p(s_0) \wedge \bigwedge_{i=0}^{k-2} T_p(s_i, s_{i+1}) \wedge T_p^-(s_{k-1}, s_k) \quad (4)$$

<sup>3</sup> PROB gives the user the opportunity to set an upper-bound on the number of successor states per event for the explicit model checker; exhaustive model checking is then not possible but counterexamples can still be found.

For the example machine in Fig. 1, we have that for  $k = 0$  the constraint remains unchanged, but for  $k = 1$  and  $k = 2$  we obtain:

$$\begin{aligned}
BMC(Inv, 0) &= m \in \{127, 255\} \wedge c_0 = 0 \wedge \neg(c_0 \leq m) \\
BMC(Inv, 1) &= m \in \{127, 255\} \wedge c_0 = 0 \\
&\quad \wedge \exists i. (i \in 1..64 \wedge c_1 = c_0 + i) \wedge c_1 \geq 0 \wedge \neg(c_1 \leq m) \\
BMC(Inv, 2) &= m \in \{127, 255\} \wedge c_0 = 0 \\
&\quad \wedge \exists i. (i \in 1..64 \wedge c_1 = c_0 + i) \wedge c_1 \geq 0 \wedge c_1 \leq m \\
&\quad \wedge \exists i. (i \in 1..64 \wedge c_2 = c_1 + i) \wedge c_2 \geq 0 \wedge \neg(c_2 \leq m)
\end{aligned}$$

Remember that  $unproven_{evt,p}(s_k)$  evaluates to true if all conjuncts of  $p$  have been proven to hold after the execution of  $evt$ . Hence, for completely proven events  $\neg unproven_{evt,p}(s_k)$  is false and the corresponding disjunct in  $T_p^-$  is obviously unsatisfiable. However, we can not remove such completely proven events from the first  $k - 1$  steps as they might contribute to the path to a violation of  $p$ , using another final event.

Another BMC approach is to use the test-case generation algorithm from [27], using  $\neg p$  as target predicate. In contrast to the BMC technique above, the transition predicate is not monolithic, and the algorithm builds up a tree of feasible paths. We have extended the algorithm from [27] to also use  $\neg unproven_{evt,p}$  instead of  $\neg p$ , where  $evt$  is the last event of any given path. The algorithm optionally uses a static enabling analysis to filter out infeasible paths before calling the solver. In the remainder of the paper we refer to this algorithm as BMC\*.

## 2.2 k-Induction

k-Induction [29] is a mixture of BMC and proof by induction. For the method to be complete, one has to avoid getting stuck in loops. Hence, the constraints are strengthened to avoid a state occurring twice on a given path.<sup>4</sup>

The base condition is encoded in Eq. (5); it is basically a BMC step and tries to find a counterexample of length  $k$  starting from the initialization. Like in Sect. 2.1 we assume that we gradually increase the value of  $k$  starting from 0, as shown in Algorithm 1. The inductive step, including the uniqueness of states, is expressed in Eq. (6), where  $Axm$  are the axioms on the constants of the model (e.g.,  $m \in \{127, 255\}$  in our running example).

$$Base(p, k) = I(s_0) \wedge \bigwedge_{i=0}^{k-1} T(s_i, s_{i+1}) \wedge \neg p(s_k) \quad (5)$$

$$Step(p, k) = Axm \wedge \bigwedge_{0 \leq i < j \leq k} s_i \neq s_j \wedge \bigwedge_{i=0}^k T(s_i, s_{i+1}) \wedge \bigwedge_{i=0}^k p(s_i) \wedge \neg p(s_{k+1}) \quad (6)$$

<sup>4</sup> We could have added these constraints  $s_i \neq s_j$  also in Sect. 2.1.

**Data:** Property  $P$   
**Result:** true iff  $P$  holds  
**1 procedure boolean** k-induction( $P$ )  
**2**      $k := 0$   
**3**     **while true do**  
**4**         **if**  $Base(P, k)$  satisfiable **then return** false  
**5**         **elseif**  $Step(P, k)$  unsatisfiable **then return** true  
**6**         **else**  $k := k + 1$  **end**  
**7**     **end**

**Algorithm 1.** k-Induction

For  $k = 0$   $Step(Inv, k)$  corresponds to trying to find counterexamples to the B invariant preservation proof obligations. In a similar fashion,  $Base(Inv, 0)$  corresponds to finding initial states which violate the invariant. Hence, if  $Base(p, 0)$  and  $Step(T, p, 0)$  are unsatisfiable, we have found an inductive proof of the property  $p$ . However, the difference with B's approach to proving invariants does appear when  $Step(T, p, 0)$  is satisfiable, i.e., there exists a state which satisfies  $p$  and a successor state violates  $p$ . The k-induction method tries to construct a real counterexample, starting from a valid initial state, not from *any* state satisfying  $p$ . Hence, the value of  $k$  is now increased and we try to find a real counterexample of length  $k + 1$  using the BMC constraint Eq. (5).

Compared to BMC, k-Induction has the advantage of including an explicit termination condition. Suppose for example we take for  $p$  the predicate  $c \geq -2 \wedge c \neq -1$  for Fig. 1. In this case BMC will never terminate, as for every value of  $k$  no counterexample can be found. k-Induction, however, can already stop with  $k = 1$ , as  $Step(Inv, 1)$  is unsatisfiable. This is an interesting result, given that the state space of the model is infinite. The constraints are shown in Fig. 4 and are unsatisfiable except for  $Step(Inv, 0)$ .  $Step(Inv, 0)$  corresponds to the B proof obligation for the event, checking that  $p$  is inductive. Hence, the B proof method is not able to prove that  $c \geq -2 \wedge c \neq -1$  always holds. (A user would need to find an inductive invariant such as  $c \geq 0$  implying the property. The IC3 algorithm in the next section will do just that automatically.)

The *Base* constraint is equal to the one in BMC:  $Base(p, k) = BMC(p, k)$ . Hence, we can include proof information in the same fashion and simply reuse the optimized constraint Eq. (4) from Sect. 2.1. For the inductive step, we can again use  $T_p^-$  for the last step, to only look for violations of *unproven* parts of  $p$ . Following Algorithm 1, we also know that all intermediate states must satisfy the property  $p$ ; this we can encode using  $T_p$ , leading to the definition in Eq. (7). As in BMC, we can not remove before-after-predicates from the first steps. Constraints are simplified but the search space is not reduced.

$$Step(p, k) = Axiom \wedge p(s_0) \wedge \bigwedge_{0 \leq i < j \leq k} s_i \neq s_j \wedge \bigwedge_{i=0}^{k-1} T_p(s_i, s_{i+1}) \wedge T_p^-(s_k, s_{k+1}) \quad (7)$$

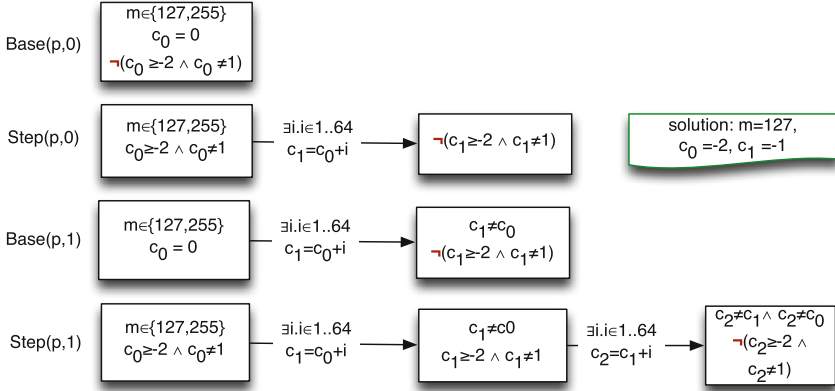


Fig. 4. Steps of k-Induction Algorithm 1 for Fig. 1 with property  $c \geq -2 \wedge c \neq -1$

### 2.3 IC3

In contrast to BMC presented in Sect. 2.1 and k-Induction presented in Sect. 2.2 the IC3 algorithm does not use an unwinding ( $\bigwedge_{i=0}^k T(s_i, s_{i+1})$ ) of the transition system. Instead, only single step queries are performed.

In order to verify a system, IC3 tries to *automatically* find an inductive invariant implying the property in question. To do so, it keeps a list of *frames*  $F_i$  over-approximating the set of states reachable in  $\leq i$  steps. Counterexamples reachable in one or two steps are handled as a special case, as shown in line 2 of Algorithm 2. Afterwards, for each level  $k$  IC3 tries to find a property violation in a single step, i. e., a solution to  $F_k \wedge T \wedge \neg p$ .

If no solution exists,  $k$  is incremented and a new frame holding  $p$  is added. Otherwise, IC3 tries to show that the faulty state is in fact not reachable from the initialization. This is done by incrementally strengthening frames until  $F_k$  becomes strong enough to prevent the property violation from occurring. A partial outline is shown in procedure *strengthen* in Algorithm 2. The *Counterexample* exception is thrown by *inductivelyGeneralize* if generalization fails and the counterexample can not be proven spurious.

Afterwards, a new counterexample might be found and IC3 will start to iterate between finding counterexamples and strengthening frames. If strengthening the frames eventually fails, a counterexample to the property is found. Otherwise, an inductive invariant has been found.

In the following, we will only go into details of IC3 wherever proof information can be incorporated. For a complete overview, see Bradley's original paper [9] or the one by Een et al. in [13]. Algorithms 2 and 3 follow the implementation of [9].

The first change to incorporate proof support takes place in the main loop of IC3. When implemented as suggested by Bradley in [9], IC3 features a special case for 0-step and 1-step reachability of a property violation as explained above.



```

Data: Property  $p$ , Transition predicate  $T$ , Initial state predicate  $I$ 
Result: true iff  $p$  holds
1 procedure boolean ic3( $p, T, I$ )
2   if  $\text{sat}(I(s_0) \wedge \neg p(s_0)) \vee \text{sat}(I(s_0) \wedge T(s_0, s_1) \wedge \neg p(s_1))$  then return false
   end
3    $F_0 := I, \text{clauses}(F_0) := \emptyset$ 
4    $F_1 := p, \text{clauses}(F_1) := \emptyset$ 
5    $k := 1$ 
6   while true do
7     if not  $\text{strengthen}(k, P, T)$  then return false end
8      $\text{propagate\_clauses}(k)$ 
9     if  $\exists i \in [1, k]: \text{clauses}(F_i) = \text{clauses}(F_{i+1})$  then return true end
10     $k := k + 1$ 
11  end

```

Algorithm 2. IC3: Main Loop

```

1 procedure boolean  $\text{strengthen}(k, p, T)$ 
2   try
3     while  $\text{sat}(F_k(s) \wedge T(s, s') \wedge \neg p(s'))$  do
4        $s :=$  the predecessor extracted from the witness
5        $n := \text{inductivelyGeneralize}(s, k - 2, k)$ 
6        $\text{pushGeneralization}((n + 1, s), k)$ 
7     end
8     return true
9   catch Counterexample
10  return false

```

Algorithm 3. IC3: Strengthen

This is shown in line 2 of Algorithm 2. The query on line 2 can be changed in the same way we did for BMC and k-Induction. After splitting the transition relation and adding proof information we obtain:

$$\text{sat}(I(s_0) \wedge \neg p(s_0)) \vee \text{sat}(I(s_0) \wedge T_p^-(s_0, s_1))$$

The key point where adding proof assistance improves the performance however is inside the *strengthen* procedure of IC3. The original version is given in Algorithm 3. Inside, the algorithm tries to find a state included in  $F_k$  that has a successor violating the property. With the usual transformation,  $F_k(s) \wedge T(s, s') \wedge \neg p(s')$  is transformed into  $F_k(s) \wedge T_p^-(s, s')$ .

In addition to simplifying the query itself, we can provide the sub routines *inductivelyGeneralize* and *pushGeneralization* with the event that lead to the violating state. This enables simplifying the respective predicates considerably.

In IC3, adding proof information has more benefits than just simplifying the occurring constraints. Due to the one-step nature of queries, constraint solving can be skipped altogether if  $\text{unproven}_{\text{evt}, p} = \text{true}$ . As no paths are built up explicitly, fully proven events have to be considered only during strengthening.

They can safely be omitted during the counterexample search. Thus, including proof information leads to a reduction of the search space.

### 3 Empirical Results

The four algorithms described above have been implemented and are available in the nightly builds of PROB<sup>5</sup>. For the empirical evaluation we want to focus on two questions:

- Does the usage of proof information considerably improve the performance of symbolic model checking algorithms for B and Event-B?
- Can symbolic model checking algorithms compete with explicit state model checking (MC) as done by PROB?

We apply both the algorithms introduced in Sect. 2 as well as PROB’s explicit state model checker (MC) to a selection of models, including artificial and real benchmarks. We use the explicit state model checker with and without proof support as outlined in [4]. The following models were used:

- *LargeBranching*, a crafted benchmark featuring a counterexample reachable in two steps. However, the initialization has numerous outgoing edges. Discovering the counterexample thus heavily relies on picking the right transitions to follow. The model is included to show that the symbolic algorithms are not influenced by this fact.
- *Search*, a classical B model of a binary search algorithm. *SearchEvents* models the same algorithm, but is written in Event-B style with simpler events. While this leads to simpler constraints, it increases the number of conjuncts due to the increased number of events.
- *TravelAgency*, a classical B model of a travel agency system storing and managing car and room rentals. The model includes an invariant violation.
- *Coloring*, a model of a graph coloring algorithm by Andriamiarina and Méry. In this particular model, the algorithm works on a concrete graph of 40 nodes.
- *f\_m0* and *f\_m1*, two hybrid models taken from [2].
- *Counters(Wrong)*, two artificial benchmarks featuring two independent counters, one of them bounded and one counting up infinitely. Both models feature an infinite state space. *CountersWrong* has a finite counterexample.
- *R0\_Gear\_Door*, *R1\_Valve*, *R2\_Outputs*, *R3\_Sensors* and *R4\_Handle* are the first refinement levels of our model [15] for the ABZ 2014 landing gear case study [8].

All benchmarks were run on a MacBook Pro featuring a 2.6 GHz i7 CPU and 8 GB of RAM. We did not run anything in parallel in order to avoid issues due to hyper-threading or scheduling. For each benchmark, a number of different results can occur:

<sup>5</sup> Available at <http://stups.hhu.de/ProB>. Information on how to use the new algorithms can be found on the PROB wiki: For the BMC\* algorithm see [http://stups.hhu.de/ProB/Bounded\\_Model\\_Checking](http://stups.hhu.de/ProB/Bounded_Model_Checking). The other algorithms are documented at [http://stups.hhu.de/ProB/Symbolic\\_Model\\_Checking](http://stups.hhu.de/ProB/Symbolic_Model_Checking).



- *verified*, i.e., the model could be model checked exhaustively without an invariant violation being detected.
- *counterexample found*, i.e., a state violating the invariant was found in the model and a trace to it has been computed.
- *incomplete*, i.e., no invariant violation has been found but model checking was not exhaustive. This could be due to timeouts or due to PROB being unable to solve occurring constraints. Currently, we do not try to recover. In case of BMC or k-Induction one could for instance try to increase  $k$  anyway.

The results are given in Table 1 showing the runtimes on successful benchmarks as well as the speedup achieved by using proof information.

The state space of the *Search* model is too large to be traversed by PROB's explicit state model checker. Unfortunately, the involved substitutions result in complex constraints that cannot be checked by the symbolic algorithms. The effect is increased by the unwinding of the transition system, as complicated constraints start to occur multiple times.

The *SearchEvents* model features simpler substitutions and is thus more suited for symbolic analysis. Using proof information, all symbolic algorithms are able to find the counterexample. Without proof information, k-Induction is not able to check the model anymore. *LargeBranching* paints a similar picture.

The *TravelAgency* model on the other hand has a relatively small state space and can easily be verified using MC. However, it features involved constructs like sequences resulting in complicated constraints. BMC\* is the only symbolic technique to find the counterexample, albeit taking much longer than MC.

The *Coloring* model is quite big and can not be checked exhaustively by MC in the given time. Only IC3 and k-Induction with proof information are able to do so. For IC3 this is due to its focus on one step reachability in combination with the model being correct: Only a small amount of counterexample candidates are discovered by IC3 and are immediately detected as spurious.

Abrial's hybrid models can be verified by MC and IC3. Here, constraints become considerably more involved with each unwinding of the transition relation done in BMC and k-Induction. IC3 is again able to verify the model thanks to its local search for counterexamples.

The infinite counters show one of the key limitations of explicit state model checking. Once a state space is infinite, exhaustive analysis is obviously impossible. For the correct model, BMC reaches its iteration limit without detecting an error. Both k-Induction and IC3 are able to analyze the models.

The landing gear model shows that the benefit of using proof information increases with the complexity of the model. As can be seen in Table 1 computation times go down once proof information is used. As for *SearchEvents* and *LargeBranching*, for *R3\_Handle* k-Induction can only successfully be used if proof information is considered. For the first refinement steps, IC3 is quicker than explicit state model checking with PROB. However, once the fourth refinement level is reached, none of the symbolic algorithms can handle the occurring constraints anymore.

Regarding speedup, we can report from  $\sim 7\%$  (*CountersWrong* with BMC) up to  $\sim 88\%$  (*R1.Valve* with k-Induction). For most of the models, incorporating proof information leads to a speedup. Using IC3, our approach leads to a performance decline for some models. We suspect it is because adding additional constraints is not necessarily beneficial for a constraint solver.

Summarizing, we can answer the two questions stated at the beginning:

- The inclusion of proof information into the symbolic model checking algorithms does improve the performance most of the time. Furthermore, some models can only be checked if proof information is used.
- For some, albeit small, models symbolic techniques can compete with explicit state model checking. Symbolic model checkers allow to verify infinite state spaces which are beyond the scope of PROB’s classical model checker.
- Among the symbolic techniques, BMC\* was the best for erroneous models, while IC3 was best for correct models.
- However, existing solvers for B and Event-B are still too weak to handle the constraints occurring in larger or more involved models. This currently hinders symbolic model checking efforts.

## 4 Discussion, Related Work and Conclusion

In [4] the authors presented a similar integration of proof information into explicit state model checking algorithms. As is the case with our implementation, the authors report a speedup by not checking invariants known to be true. In contrast to our approach, the use of proof information never slowed down the model checking process.

Compared with [4], we have added a way to construct proof information within PROB itself, using a bridge to the Atelier B provers and using PROB’s proving capabilities [18]. Of course this takes time and does not always pay off.

In [4], as with BMC and k-Induction, the search space itself is never reduced. Search space reduction through using proof techniques is considered in [28] and [24]. For model checking CTL and LTL properties, proof information can be used as well. In [26] the model checker SMV is coupled with theorem proving techniques. In a similar fashion, [3] combines the Alloy Analyzer with the Athena theorem prover.

Instead of using theorem provers to support model checking, one can use model checkers for theorem proving. We have done so using PROB [18,22].

Our evaluation shows that using symbolic model checking techniques for B and Event-B models is beneficial: Several counterexamples could only be detected by the symbolic algorithms. Furthermore, some models could be model checked exhaustively. As already outlined in [14], symbolic techniques prove to be a valuable addition to explicit techniques. The techniques are actually also applicable to  $TLA^+$ , via PROB’s translation from  $TLA^+$  to B [16]<sup>6</sup>.

<sup>6</sup> For further information regarding  $TLA^+$  support in ProB have a look at <http://stups.hhu.de/ProB/TLA>.

The key weakness of employing symbolic model checking techniques lies within the expressiveness of B and Event-B. Even though constraint solvers and SMT solvers have increased their efficiency by a huge margin, the constraints occurring during symbolic model checking of high-level languages like B are still too involved. Among other abstraction techniques, integrating static (proof) information into the constraints is one way to help. It brings down computation times and sometimes enables successful validation. We are also working on strengthening the underlying constraint solver, by integrating SMT solvers such as Z3 [19]. Still, more improvements need to be achieved until full symbolic verification of B and Event-B models becomes viable.

Regarding the different model checking algorithms, especially IC3 seems promising. In contrast to the other two algorithms, its focus on one step reachability keeps occurring constraints easier. This makes it more suited for symbolic model checking of high-level languages like B and Event-B. Additionally, the integration of proof information can lead to a reduced search space. As IC3 has originally been developed for hardware model checking, it is not trivial to lift it to the software world. To do so, we would like to investigate IC3 for B together with abstraction techniques as introduced in [11] or [7].

Another direction of future work could be to generate missing proof obligations from the model checking run. Analyzing predicates that lead to a timeout one could find problematic properties and try to prove them externally or in an independent run. Once the constraint solver gets stuck we could ask an external solver<sup>7</sup> to prove or disprove further invariants. Afterwards, one could extend the set of properties under consideration.

In summary, we have implemented four symbolic model checking algorithms for B and Event-B and have shown how to integrate proof information to improve the algorithms' performance. Our evaluation shows that bounded model checking can effectively find counterexamples in models with very large branching factors and that IC3 is capable of automatically proving models with infinite state spaces correct. Further research is, however, needed to scale up the symbolic techniques to models with more involved events.

**Acknowledgements.** We would like to thank the reviewers of ABZ'2016 for their useful feedback. We also thank Aymerick Savary for comments and ideas, in particular relating to BMC and test-case generation.

## References

1. Abrial, J.-R., Butler, M., Hallerstede, S., Hoang, T., Mehta, F., Voisin, L.: Rodin: an open toolset for modelling and reasoning in Event-B. *Int. J. Softw. Tools Technol. Transf.* **12**(6), 447–466 (2010)
2. Abrial, J.-R., Su, W., Zhu, H.: Formalizing hybrid systems with event-B. In: Derrick, J., Fitzgerald, J., Gnesi, S., Khurshid, S., Leuschel, M., Reeves, S., Riccobene, E. (eds.) ABZ 2012. LNCS, vol. 7316, pp. 178–193. Springer, Heidelberg (2012)

<sup>7</sup> Like the Atelier B provers or the SMT solvers for Rodin.

3. Arkoudas, K., Khurshid, S., Marinov, D., Rinard, M.: Integrating model checking and theorem proving for relational reasoning. In: Berghammer, R., Möller, B., Struth, G. (eds.) *ReLMiCS 2003*. LNCS, vol. 3051, pp. 21–33. Springer, Heidelberg (2004)
4. Bendisposto, J., Leuschel, M.: Proof assisted model checking for B. In: Breitman, K., Cavalcanti, A. (eds.) *ICFEM 2009*. LNCS, vol. 5885, pp. 504–520. Springer, Heidelberg (2009)
5. Biere, A.: Bounded model checking. In: *Handbook of Satisfiability*, pp. 457–481 (2009)
6. Biere, A., Cimatti, A., Clarke, E., Zhu, Y.: Symbolic model checking without BDDs. In: Cleaveland, W.R. (ed.) *TACAS 1999*. LNCS, vol. 1579, p. 193. Springer, Heidelberg (1999)
7. Birgmeier, J., Bradley, A.R., Weissenbacher, G.: Counterexample to induction-guided abstraction-refinement (CTIGAR). In: Biere, A., Bloem, R. (eds.) *CAV 2014*. LNCS, vol. 8559, pp. 831–848. Springer, Heidelberg (2014)
8. Boniol, F., Wiels, V.: The landing gear system case study. In: Boniol, F., Wiels, V., Ait Ameer, Y., Schewe, K.-D. (eds.) *ABZ 2014*. CCIS, vol. 433, pp. 1–18. Springer, Heidelberg (2014)
9. Bradley, A.R.: SAT-based model checking without unrolling. In: Jhala, R., Schmidt, D. (eds.) *VMCAI 2011*. LNCS, vol. 6538, pp. 70–87. Springer, Heidelberg (2011)
10. Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.J.: Symbolic model checking:  $10^{20}$  states and beyond. *Inf. Comput.* **98**(2), 142–170 (1992)
11. Cimatti, A., Griggio, A.: Software model checking via IC3. In: Madhusudan, P., Seshia, S.A. (eds.) *CAV 2012*. LNCS, vol. 7358, pp. 277–293. Springer, Heidelberg (2012)
12. Déharbe, D., Fontaine, P., Guyot, Y., Voisin, L.: Integrating SMT solvers in Rodin. *Sci. Comput. Program.* **94**(P2), 130–143 (2014)
13. Een, N., Mishchenko, A., Brayton, R.: Efficient implementation of property directed reachability. In: *Proceedings of the International Conference on Formal Methods in Computer-Aided Design (FMCAD 2011)*, pp. 125–134, Austin, TX, FMCAD Inc (2011)
14. Hallerstede, S., Leuschel, M.: Constraint-based deadlock checking of high-level specifications. *Theor. Pract. Logic Program.* **11**(4–5), 767–782 (2011)
15. Hansen, D., Ladenberger, L., Wiegard, H., Bendisposto, J., Leuschel, M.: Validation of the ABZ landing gear system using ProB. In: Boniol, F., Wiels, V., Ait Ameer, Y., Schewe, K.-D. (eds.) *ABZ 2014*. CCIS, vol. 433, pp. 66–79. Springer, Heidelberg (2014)
16. Hansen, D., Leuschel, M.: Translating  $TLA^+$  to B for validation with ProB. In: Derrick, J., Gnesi, S., Latella, D., Treharne, H. (eds.) *IFM 2012*. LNCS, vol. 7321, pp. 24–38. Springer, Heidelberg (2012)
17. Hansen, D., Leuschel, M.: Translating B to  $TLA^+$  for validation with TLC. In: Ait Ameer, Y., Schewe, K.-D. (eds.) *ABZ 2014*. LNCS, vol. 8477, pp. 40–55. Springer, Heidelberg (2014)
18. Krings, S., Bendisposto, J., Leuschel, M.: From failure to proof: the prob disprover for B and event-B. In: Calinescu, R., Rumpe, B. (eds.) *SEFM 2015*. LNCS, vol. 9276, pp. 199–214. Springer, Heidelberg (2015)
19. Krings, S., Leuschel, M.: SMT Solvers for Validation of B and Event-B models. In: *Proceedings iFM’2016*, LNCS. Springer (2016). to appear

20. Leuschel, M., Butler, M.: ProB: a model checker for B. In: Araki, K., Gnesi, S., Mandrioli, D. (eds.) FME 2003. LNCS, vol. 2805, pp. 855–874. Springer, Heidelberg (2003)
21. Leuschel, M., Butler, M.: ProB: an automated analysis toolset for the B method. *Int. J. Softw. Tools Technol. Transf.* **10**(2), 185–203 (2008)
22. Ligot, O., Bendisposto, J., Leuschel, M.: Debugging event-B models using the ProB disprover plug-in. In: Proceedings AFADL 2007, June 2007
23. Matos, P.J., Fischer, B., Marques-Silva, J.: A lazy unbounded model checker for EVENT-B. In: Breitman, K., Cavalcanti, A. (eds.) ICFEM 2009. LNCS, vol. 5885, pp. 485–503. Springer, Heidelberg (2009)
24. Müller, O., Nipkow, T.: Combining model checking and deduction for I/O-automata. In: Brinksma, E., Steffen, B., Cleaveland, W.R., Larsen, K.G., Margaria, T. (eds.) TACAS 1995. LNCS, vol. 1019, pp. 1–16. Springer, Heidelberg (1995)
25. Plagge, D., Leuschel, M.: Validating B, Z and TLA<sup>+</sup> using PROB and Kodkod. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012. LNCS, vol. 7436, pp. 372–386. Springer, Heidelberg (2012)
26. Pnueli, A., Ruah, S., Zuck, L.D.: Automatic deductive verification with invisible invariants. In: Margaria, T., Yi, W. (eds.) TACAS 2001. LNCS, vol. 2031, p. 82. Springer, Heidelberg (2001)
27. Savary, A., Frappier, M., Leuschel, M., Lanet, J.-L.: Model-based robustness testing in event-B using mutation. In: Calinescu, R., Rumpe, B. (eds.) SEFM 2015. LNCS, vol. 9276, pp. 132–147. Springer, Heidelberg (2015)
28. Shankar, N.: Combining theorem proving and model checking through symbolic analysis. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, p. 1. Springer, Heidelberg (2000)
29. Sheeran, M., Singh, S., Stålmarch, G.: Checking safety properties using induction and a SAT-solver. In: Johnson, S.D., Hunt Jr., W.A. (eds.) FMCAD 2000. LNCS, vol. 1954, pp. 108–125. Springer, Heidelberg (2000)
30. Witulski, J., Leuschel, M.: Checking computations of formal method tools - a secondary toolchain for prob. In: Proceedings of the 1st Workshop on Formal-IDE (EPTCS), *Electronic Proceedings in Theoretical Computer Science*, vol. 149 (2014)
31. Yu, Y., Manolios, P., Lampert, L.: Model checking TLA<sup>+</sup> specifications. In: Pierre, L., Kropf, T. (eds.) CHARME 1999. LNCS, vol. 1703, pp. 54–66. Springer, Heidelberg (1999)