

Abstraction Refinement in Symbolic Model Checking Using Satisfiability as the Only Decision Procedure*

Bing Li, Chao Wang, Fabio Somenzi

University of Colorado at Boulder
e-mail: {bli, wangc, Fabio}@Colorado.EDU

Received: date / Revised version: date

Abstract We present an abstraction refinement algorithm for model checking of safety properties that relies exclusively on a SAT solver for checking the abstract model, testing abstract counterexamples on the concrete model, and refinement. Model checking of the abstractions is based on bounded model checking extended with checks for the existence of simple paths that help in deciding passing properties. All minimum-length spurious counterexamples are eliminated in one refinement step by an incremental procedure that combines the analysis of the conflict dependency graph produced by the SAT solver while looking for concrete counterexamples with an effective refinement minimization procedure.

1 Introduction

Model checking [CGP99] is an algorithmic approach to the verification of properties of reactive systems, which has been successfully applied to both hardware and software. Since model checking entails the exploration of a potentially very large state space, the alleviation of the so-called state explosion problem has been the object of much research. On the one hand, techniques have been developed that allow models with hundreds of state variables to be analyzed directly. On the other hand, abstraction has been used to allow the model checker to draw conclusions on the original, concrete model by examining a simpler, abstract one.

For systems with many state variables and many transitions, the symbolic approach has proved crucial. In symbolic model checking, sets of states and transition are described by their characteristic functions. Various forms of representation have been used for these functions, the most popular being Binary Decision Diagrams (BDDs) [Bry86], and Conjunctive Normal Form (CNF).

Classical BDD-based model checking [McM94] is based on the computation of fixpoints. For instance, the reachable states of a model are computed as the least fixpoint of the

function $\lambda Z. I \vee \text{Succ}(Z)$, which adds the successors of the states in Z to the initial states. Both the set of states and the successor relation are stored as BDDs. The fixpoint computation converges in a number of iterations that equals the maximum distance of a reachable state from the initial states. Checking for convergence is made easy by the strong canonicity of BDDs (identical sets share the same representation). BDD-based model checking can therefore prove properties almost as easily as it can disprove them.

Bounded Model Checking (BMC) [BCCZ99], on the other hand, formulates the reachability test as a series of satisfiability (SAT) checks for paths of bounded length. (To see if a path of length k to a set of states exists, the transition relation is unrolled k times.) For finite systems the process must eventually terminate: The length of the shortest path between two states cannot exceed the number of states. Hence, if no path is found with length up to the number of states, the target states are known to be unreachable. This observation, however, does not help for the kind of models that one encounters in practice. The diameter of the state graph would give a much better bound on k , but, unfortunately, it is hard to compute [BCCZ99]. For this reason, BMC has come to be regarded as an excellent *debugging* (as opposed to *verification*) technique. That is, classical BMC is particularly adept at finding counterexamples, but ill-suited to prove their absence.

The ability demonstrated by BMC to deal with models beyond the reach of BDD-based methods has sparked interest in the use of CNF and SAT for proof as well as refutation. Two main approaches have been pursued: The replacement of BDDs with CNF formulae in the fixpoint computation [ABE00, WBCG00, McM02], and the development of more effective termination criteria for BMC.

The opportunity of replacing BDDs with CNF formulae can be argued on the grounds that canonicity of representation makes BDDs somewhat inflexible. Hence, some functions that admit compact representations in CNF have exceedingly large BDDs. However, the inflexibility argument can also be used against CNF, and memoization techniques are more effective for BDDs. In fact, to date, CNF-based fix-

* This work was supported in part by SRC contract 2003-TJ-920.

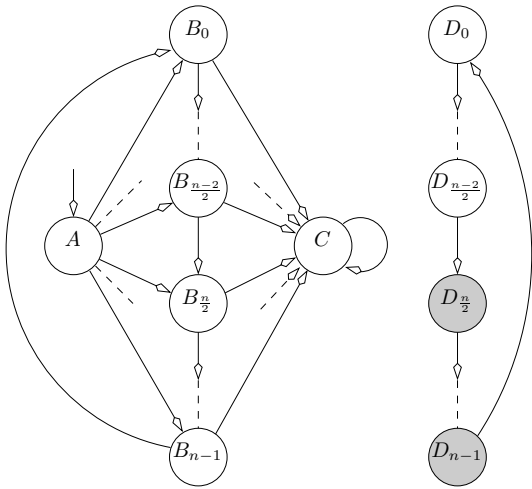


Fig. 1 Model with long simple path

point computation has not demonstrated a consistent advantage over the classical BDD-based one. One may argue that the main reason for the success of BMC in finding counterexamples lies in its avoidance of the needless computation and storage of reachable states that are not on the error trace.

Several proposals have been made to improve BMC’s ability to prove the non-existence of a path. It is straightforward to check for inductive invariants, since it only entails checking for the existence of a transition from a state that satisfies the invariant to one that does not. An extension of the inductive approach has been presented in [SSS00], in which termination occurs as soon as the length of the path reaches the length of the longest simple path from an initial state, or to a target state. A recent paper [McM03] proposes the analysis of the unsatisfiable formulae to allow termination when the *reverse sequential depth* of the model is reached.

Early termination in BMC requires additional checks beyond the one for the existence of paths of certain lengths. These checks translate into more clauses in the CNF formulae whose satisfiability has to be established. For the approach of [SSS00], the number of extra clauses is quadratic in the length of the path. As a result, it is not surprising that finding counterexamples is slower than with pure BMC. The extra cost, however, appears to be worth paying, since it increases substantially the fraction of passing properties that can be decided. Unfortunately, there remain instances for which the additional termination tests are too expensive¹. Consider the model illustrated in Fig. 1. It has $2n+2$ states, one of which is initial (A). The $n/2$ states $D_{n/2}, \dots, D_{n-1}$ are the (unreachable) target states. The longest simple path from the initial state has length $n+1$, while the longest simple path to a target state that does not visit any other target state has length $n/2$; the reverse sequential depth of the model is also $n/2$. Hence, the methods of [SSS00, McM03] will have to consider

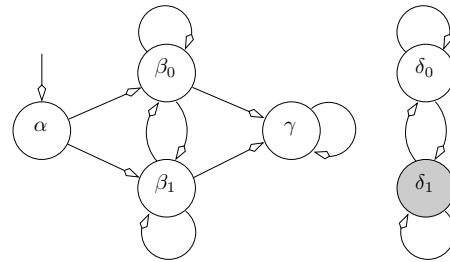


Fig. 2 Abstraction of the model of Fig. 1

paths of length $n/2$ before they can declare the target states unreachable. By contrast, the forward sequential depth is 2.

Fig. 2 shows an abstraction of the model of Fig. 1. States A, B_i, C , and D_i are abstracted by $\alpha, \beta_{\lfloor 2i/n \rfloor}, \gamma$, and $\delta_{\lfloor 2i/n \rfloor}$, respectively. The target state remains unreachable in this model, and the forward sequential depth is still 2; however, the longest simple path and the sequential depth are reduced. Though in general there is no guarantee that abstraction will shorten or even not lengthen the longest simple paths, or the shortest paths, this example illustrates how abstraction may help BMC, especially for passing properties.

Abstraction and BMC have been combined in more than one recent work, especially in the context of abstraction refinement. In abstraction refinement [Kur94], one starts with a coarse abstraction of the given, concrete model and keeps refining it until the property is decided. For universal properties like the reachability properties that are the focus of this article, this often means that the abstract models simulate the concrete one [Mil71], and that either the property is shown to hold on an abstract model, or a counterexample is found in the concrete one. In [WHL⁺01, CGKS02, CCK⁺02, WLJ⁺03] BMC is used to check whether counterexamples found in the abstract models can be *concretized*, that is, whether a counterexample can be found in the concrete model that is mapped by the abstraction onto the abstract counterexample. The first three of these methods also analyze the failed concretization test to guide the refinement. Therefore, they represent instances of counterexample-guided abstraction refinement. On the other hand, [WLJ⁺03] analyzes the abstract model to decide how to refine it. Yet another approach is the one of [MA03], in which the abstract model is derived from a failing BMC run on the concrete model. This reversal of the customary order is attractive for those frequent cases in which paths of moderate length can be easily checked on the concrete model.

One common trait of the approaches to abstraction refinement mentioned so far is the application of a BDD-based model checker to the abstract models, and of SAT solvers to the concrete ones. By contrast, the objective of this article is to explore what can be achieved with a SAT solver as the only decision procedure in the abstraction refinement framework. The rationale for combining BDDs and SAT is that each is well-suited to the task assigned to it: The SAT solver is good at checking the existence of a path of a given length in a large model, whereas the BDD-based model checker is bet-

¹ In [KS03] it was shown how to reduce the number of clauses to $O(k \log^2 k)$. However, the extra clauses to enforce simple paths still represent a significant hurdle for the SAT solvers.

ter at proving the absence of certain paths, regardless of their lengths, in a model of moderate size. This observation is certainly well motivated when one regards the models for which abstraction refinement results have been reported in the literature; their sizes rarely exceed 1,000 binary state variables. As the models grow larger, however, we expect an approach purely based on SAT to become more competitive. Therefore, our goal is to eventually be able to switch between BDD-based model checking and SAT-based techniques for the analysis of the concrete model. In this article we report on a significant step in that direction by presenting an algorithm for abstraction refinement that is purely based on SAT.

Our approach is similar to the ones discussed so far in the fact that abstractions are obtained by removing part of the state variables of the model; refinement then consists of reinstating some of the removed variables. The algorithm has four major components: the decision procedure for the abstract model is the one of [SSS00], which has already been mentioned. The second component is the concretization test: as in [BGG02], a gradual refinement approach tries to prove counterexamples spurious without resort to the concrete model. The third component—the choice of the refinement—makes use of elements of [WLJ⁺03], [MA03], and [CCK⁺02]. Like the first two, it addresses all the abstract counterexamples of a certain length at once; like the last two, it analyzes the proof of non-existence of counterexamples of a certain length to derive a set of candidate state variables from which the ones that will be added to the abstract model are chosen. Finally, the fourth component is a heuristic procedure for abstraction minimization. This minimization is quite important in our approach, because the simultaneous elimination of all spurious counterexamples of a certain length tends to generate large sets of candidate variables. Our experimental evaluation of the SAT-based abstraction refinement algorithm compared it to both BMC (with and without termination checks for passing properties) and to the best abstraction refinement algorithm available to us [WLJ⁺03]. The results, discussed in Section 5, show that the new approach, though not uniformly superior, is more robust than the others, and is especially promising for the more challenging problems.

2 Preliminaries

2.1 Open Systems

Let $V = \{v_1, \dots, v_n\}$ and $W = \{w_1, \dots, w_m\}$ be sets of Boolean variables. We designate by V' the set $\{v'_1, \dots, v'_n\}$ consisting of the primed version of the elements of V , and by V^i the set $\{v_1^i, \dots, v_n^i\}$. Likewise, $W^i = \{w_1^i, \dots, w_m^i\}$. An *open system* is a 4-tuple

$$\langle V, W, I, T \rangle,$$

where V is the set of (current) state variables, W is the set of combinational variables, $I(V)$ is the initial state predicate, and $T(V, W, V')$ is the transition relation. The variables in V'

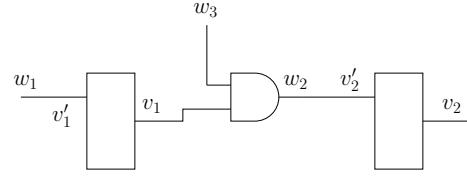


Fig. 3 A sequential circuit defining an open system with $n = 2$ and $m = 3$. The rectangles represent binary state elements

are the next state variables. All sets are finite, and all variables range over finite domains.

We assume that $T(V, W, V')$ is given by a *circuit graph*, that is, by a labeled graph $\mathcal{C} = (V \cup W, E)$ such that $m \geq n$, node $v_i \in V$ is labeled by $w_i \in W$, node $w_i \in W$ is labeled by a Boolean formula $T_i = w_i \leftrightarrow \delta_i(V, W)$, $(w_i, v_i) \in E$ for $i \in \{1, \dots, n\}$, and, for $x \in V \cup W$, $w_i \in W$, $(x, w_i) \in E$ iff x appears in δ_i . The transition relation is then defined by:

$$T(V, W, V') = \bigwedge_{1 \leq i \leq n} (v'_i \leftrightarrow w_i) \wedge \bigwedge_{1 \leq i \leq m} T_i(W, V). \quad (1)$$

In a sequential circuit, the variables in W are associated with the primary inputs and the outputs of the combinational logic gates of the circuit; the variables in V are associated with the memory elements. Each T_i is called a *gate relation* because it usually describes the behavior of a logic gate. For instance, if w_i is the output variable of a two-input AND gate with inputs w_j and v_k , then $T_i = w_i \leftrightarrow (w_j \wedge v_k)$. If, on the other hand, w_i is a primary input to the circuit, then $T_i = 1$. Each term of the form $v'_i \leftrightarrow w_i$ equates a next state variable to a combinational variable. (The output of the gate feeding the i -th memory element.)

A state variable v_j is said to be in the *direct support* of variable v_i (w_i) if v_j is connected to v_i (w_i) by a path in \mathcal{C} that goes through nodes in W (logic gates) only. Variable v_j is in the *cone of influence* (COI) of v_i (w_i) if there is a path (of any kind) connecting v_j to v_i (w_i) in \mathcal{C} .

Example 1 Figure 3 shows a simple sequential circuit with two state variables. The transition relation (1) is given by

$$T(V, W, V') = (v'_1 \leftrightarrow w_1) \wedge (v'_2 \leftrightarrow w_2) \wedge (w_2 \leftrightarrow (w_3 \wedge v_1)).$$

Note that $T_1 = T_3 = \text{true}$. If the only initial state of the circuit is $v_1 = 0, v_2 = 0$, then $I(V) = \neg v_1 \wedge \neg v_2$. Variable v_1 is in the direct support of w_2 , and in the COI of v_2 . \square

2.2 Proving Safety Properties

An open system Ω defines a labeled transition structure in the usual way, with states Q_Ω corresponding to the valuations of the variables in V , and transition labels corresponding to the valuations of the variables in W . Conversely, a set of states $S \subseteq Q_\Omega$ corresponds to a predicate $S(V)$ or $S(V')$. Predicate $S(V)$ ($S(V')$) is the characteristic function of S expressed in terms of the current (next) state variables. State $q \in Q_\Omega$ is an initial state if it satisfies $I(V)$. State set $S \subseteq Q_\Omega$ is *reachable*

from state set S' in k steps if there is a path of length k in the labeled transition structure defined by Ω that connects some state in S' to some state in S ; equivalently if

$$S'(V^0) \wedge \bigwedge_{1 \leq i \leq k} T(V^{i-1}, W^i, V^i) \wedge S(V^k) \quad (2)$$

is satisfiable. State set S is reachable from S' if there exists $k \in \mathbb{N}$ such that S is reachable in k steps from S' . A state set is reachable (in k steps) if it is reachable (in k steps) from I . When no confusion arises we shall identify a state $q \in Q_\Omega$ with the set $\{q\}$. A finite (infinite) sequence of states $\rho \in Q_\Omega^*$ ($\in Q_\Omega^\omega$) is a finite (infinite) *run* of Ω if the first state is initial, and every other state is reachable from its predecessor in one step. The set of all possible runs of Ω is the language of Ω , denoted by $\mathcal{L}(\Omega)$.

A linear-time *safety property* P of Ω is a subset of Q_Ω^ω such that any infinite sequence over Q_Ω not in P has a finite prefix that cannot be extended to a sequence in P [AS85]. Open system Ω satisfies safety property P if $\mathcal{L}(\Omega) \subseteq P$. Checking the satisfaction of an ω -regular safety property P by an open system Ω can be reduced to the reachability problem by composing Ω with an automaton \mathcal{A}_P that accepts the inextensible prefixes of the sequences not in P . The property is satisfied by the open system if no state of the composition $\Omega \parallel \mathcal{A}_P$ that projects on an accepting state of \mathcal{A}_P is reachable. In the sequel we restrict ourselves to ω -regular safety properties, and assume that the given open system already incorporates the property automaton. This assumption allows us to identify the property with a set of states of the system, which we also denote by P . Hence, property P is satisfied by Ω if there is no $k \in \mathbb{N}$ such that

$$I(V^0) \wedge \bigwedge_{1 \leq i \leq k} T(V^{i-1}, W^i, V^i) \wedge \neg P(V^k) \quad (3)$$

is satisfiable. An invariant is a safety property that states that a certain predicate holds of all reachable states of Ω . In this case P is the set of states that satisfy that predicate.

Example 2 Continuing Example 1, suppose that $P(V) = \neg v_2$. Then a counterexample of length 2 to P can be found by solving

$$\begin{aligned} & \neg v_1^0 \wedge \neg v_2^0 \wedge (v_1^1 \leftrightarrow w_1^1) \wedge (v_2^1 \leftrightarrow w_2^1) \wedge \\ & (w_2^1 \leftrightarrow (w_3^1 \wedge v_1^0)) \wedge (v_1^2 \leftrightarrow w_1^2) \wedge \\ & (v_2^2 \leftrightarrow w_2^2) \wedge (w_2^2 \leftrightarrow (w_3^2 \wedge v_1^1)) \wedge v_2^2 . \end{aligned}$$

A satisfying assignment for this formula is

$$\neg v_1^0 \wedge \neg v_2^0 \wedge w_1^1 \wedge \neg w_2^1 \wedge v_1^1 \wedge \neg v_2^1 \wedge w_1^2 \wedge w_2^2 \wedge w_3^2 \wedge v_1^2 \wedge v_2^2 .$$

The value of w_3^1 is irrelevant. \square

The search for a k such that (3) is satisfiable can obviously be restricted to the range $\{0, \dots, |Q_\Omega| - 1\}$. Hence, in theory, the process is guaranteed to terminate. In practice, the number of states is too large to be of any practical use, and tighter upper bounds for k are sought. In model checking

approaches that are based on fixpoint computations [McM94, ABE00, WBCG00, McM02], the maximum value of k is provided by the number of iterations needed to reach convergence. On the other hand, for algorithms that directly check the satisfiability of (3), the diameter of the graph [BCCZ99] or bounds obtained from the structure of the hardware model have been proposed [BKA02]. Here we summarize a method proposed in [SSS00] that is of particular interest to us.

A *simple path* is one that visits a state at most once. If some state in $\neg P$ is reachable, there must exist a simple path from an initial state to it that does not go through any other states in I or $\neg P$. Hence, if no simple path of length k exists such that its first state is initial and no other state is initial, or such that its final state is in $\neg P$ and no other state is in $\neg P$, then, there is no path of length greater than or equal to k connecting a state in I to a state in $\neg P$. If in addition, there is no path of length less than k connecting I to $\neg P$, then $\Omega \models P$. Two sets of states S' and S are connected by a simple path of length k in Ω if

$$\begin{aligned} \Sigma_k(S', S) = & S'(V^0) \wedge \bigwedge_{1 \leq i \leq k} T(V^{i-1}, W^i, V^i) \\ & \wedge S(V^k) \wedge \bigwedge_{0 \leq j < i \leq k} \bigvee_{1 \leq l \leq n} (v_l^i \neq v_l^j) \quad (4) \end{aligned}$$

is satisfiable. Checking the two conditions above then amounts to checking that either of the following formulae is unsatisfiable:

$$\Sigma_k(I, Q) \wedge \bigwedge_{0 < i \leq k} \neg I(V^i) \quad (5)$$

$$\Sigma_k(Q, \neg P) \wedge \bigwedge_{0 \leq i < k} P(V^i) . \quad (6)$$

Note that the predicate corresponding to the set Q is true.

2.3 Abstraction Refinement

Abstract interpretation [CC77] provides a very flexible framework for the description of abstraction. In this article, however, we consider the following restricted definition. Open system $\hat{\Omega} = \langle \hat{V}, \hat{W}, \hat{I}, \hat{T} \rangle$ is an *abstraction* of Ω if

$$\begin{aligned} & - \hat{V} \subseteq V; \\ & - \hat{W} = \hat{W}_1 \cup (V \setminus \hat{V}); \\ & - \hat{W}_1 \subseteq W \text{ such that } v_i \in \hat{V} \text{ implies } w_i \in \hat{W}; \\ & - \hat{I}(\hat{V}) = \exists(V \setminus \hat{V}) . I(V); \\ & - \hat{T}(\hat{V}, \hat{W}, \hat{V}') = \exists(V \setminus \hat{V}) . \exists(W \setminus \hat{W}) . \\ & \quad \exists(V' \setminus \hat{V}') . T(V, W, V'). \end{aligned}$$

(Note that w_i is the combinational variable associated to v'_i .) This definition entails that $\hat{\Omega}$ simulates [Mil71] Ω . Hence, every run of Ω has a matching run in $\hat{\Omega}$. Property \hat{P} is the abstraction of property P with respect to $\hat{\Omega}$ if $\hat{P}(\hat{V}) = \forall(V \setminus \hat{V}) . P(V)$. If P is an ω -regular property and $\hat{\Omega}$ satisfies (or models) \hat{P} , then Ω satisfies P . That is,

$$\hat{\Omega} \models \hat{P} \rightarrow \Omega \models P . \quad (7)$$

This preservation result is the basis for the following abstraction refinement approach to the verification of P . One starts with a coarse abstraction $\widehat{\Omega}_0$ of the *concrete* open system Ω and checks whether $\widehat{\Omega}_0 \models \widehat{P}_0$. If that is the case, then $\Omega \models P$; otherwise, there exists a least $k' \in \mathbb{N}$ such that

$$\widehat{I}(\widehat{V}^0) \wedge \bigwedge_{1 \leq i \leq k'} T(\widehat{V}^{i-1}, \widehat{W}^i, \widehat{V}^i) \wedge \neg \widehat{P}(\widehat{V}^{k'}) \quad (8)$$

is satisfiable. The satisfying assignments to (8) are the shortest-length *abstract counterexamples* (ACEs). If $\widehat{\Omega}_0 \not\models \widehat{P}_0$ one or more ACEs are checked for *concretization*. That is, one checks whether (3) has solutions that agree with the ACE(s) being checked. Because of the additional constraints provided by the ACEs, a concretization test is often less expensive than the satisfiability check of (3). However, its failure only indicates that the abstract error traces are spurious. Therefore, if the concretization test fails, one chooses a refined abstraction $\widehat{\Omega}_1$ and repeats the process, until one of these cases occurs.

1. $\widehat{\Omega}_i \models \widehat{P}_i$ for some i , in which case $\Omega \models P$ is inferred.
2. The concretization test passes for some i , in which case it is concluded that $\Omega \not\models P$ and the satisfying assignment found is returned as counterexample to P .
3. The refinement eventually produces $\widehat{\Omega}_i = \Omega$. In this final case, the satisfiability check of (8) answers the model checking question conclusively. This is an undesirable outcome because the purpose of abstraction is defeated.

When the refinement $\widehat{\Omega}_{i+1}$ of $\widehat{\Omega}_i$ is chosen with the help of the information provided by the failed concretization test, one talks of counterexample-guided abstraction refinement.

The cone of influence (direct support) of a property is the union of the cones of influence (direct supports) of all the variables mentioned in the predicate that defines the property. Cone-of-influence reduction refers to the abstraction in which \widehat{V} is the COI of the property. It is commonly applied before any model checking is attempted, because it satisfies

$$\widehat{\Omega} \models \widehat{P} \leftrightarrow \Omega \models P . \quad (9)$$

Example 3 For the circuit of Fig. 3, consider the following abstraction $\widehat{\Omega}$:

- $\widehat{V} = \{v_2\}$;
- $\widehat{W} = \{w_2, w_3, v_1\}$;
- $\widehat{I}(\widehat{V}) = \neg v_2$;
- $\widehat{T}(\widehat{V}, \widehat{W}, \widehat{V}') = (v_2' \leftrightarrow w_2) \wedge (w_2 \leftrightarrow (w_3 \wedge v_1))$.

Suppose $P(V) = \widehat{P}(\widehat{V}) = \neg v_2$. Then from (8) one finds that a counterexample of length 1 exists in $\widehat{\Omega}$ by solving

$$\neg v_2^0 \wedge (v_2^1 \leftrightarrow w_2^1) \wedge (w_2^1 \leftrightarrow (w_3^1 \wedge v_1^1)) \wedge v_2^1 .$$

The only satisfying assignment is $\neg v_2^0 \wedge w_2^1 \wedge w_3^1 \wedge v_1^1 \wedge v_2^1$. Note that this ACE cannot be concretized because v_1^0 cannot be set to 1 in the concrete model. After refinement adds v_1 to \widehat{V} , $\widehat{\Omega} = \Omega$. The counterexample found in Example 2 therefore shows that $\Omega \not\models P$. \square

2.4 Satisfiability Solvers

Many modern SAT solvers are based on clause recording. Whenever they detect a conflicting assignment to a formula f (one that causes f to evaluate to false), they conjoin a *conflict clause* to f . The new clause prevents the solver from attempting the same assignment again. It may also exclude from future consideration other parts of the search space that can be inferred to contain no satisfying assignments. Such inference is based on the analysis of the so-called *implication graph*, which shows which decisions and clauses are responsible for the conflict. A SAT solver maintains the implication graph by recording, for each assignment, the *level* at which it was made², whether it was a decision, and if not, the clause that implied it. When a conflict is detected, a cut separating the sink of the implication graph from the sources identifies a set of assignments sufficient to cause the conflict. The disjunction of the negation of those assignments is a conflict clause.

The clauses that make up the edges of an implication graph between the cut and the sink can be used to explain the conflict clause deduced from it. When using a SAT solver in model checking based on abstraction refinement, there are two important reasons to keep track of the explanations of conflict clauses. The first reason is the ability to identify an *unsatisfiable core* [GN03,ZM03] of the given formula when it is indeed unsatisfiable. This can be done by recursively replacing conflict clauses with those appearing between the chosen cut and the sink in the implication graphs that produced them. The process starts from the final conflict clause and terminates when only clauses of the original formula are left. When the check for existence of counterexamples of a certain length fails, the unsatisfiable core produced by the SAT solver can be used to guide the refinement of the abstract model.

The second reason to preserve the mapping between conflict clauses and the edges of their implication graphs is to help in the incremental solution of sequences of SAT instances [WKS01,ES03]. If f and f' are two CNF formulae, and γ is a conflict clause of f , then γ is also a conflict clause of f' if f' contains all the clauses of f in the implication graph for γ up to the cut that identifies γ . Therefore, if f' is obtained by slight modification of f , many conflict clauses derived in checking the satisfiability of f may be added to f' inexpensively if their “explanations” in terms of clauses of f are saved. As a special case, if f' contains all the clauses of f , then all conflict clauses deduced while checking the satisfiability of f are also valid for f' . Sequences of closely related SAT instances are found in our algorithm for abstraction refinement. They are discussed in Section 3.4 together with the exploitation of inherited conflict clauses to speed up SAT checking.

² The level starts from 0; it is incremented at every new decision; and it decreases when the algorithm backtracks.

```

boolean PURESAT( $\Omega, P, \mathcal{C}$ ) {
1   $L = 0$ ;
2   $\hat{\Omega} = \text{CREATEINITIALABSTRACTION}(\Omega, P)$ ;
3  while ( $\hat{\Omega} \neq \Omega$ ) {
4    if ( $\neg \text{CHECKSIMPLEPATH}(\hat{\Omega}, P, L)$ )
5      return TRUE;
6     $\tilde{\Omega} = \hat{\Omega}$ ;
7    while ( $\text{EXISTCEX}(\tilde{\Omega}, P, L)$ ) {
8      if ( $\tilde{\Omega} == \Omega$ )
9        return FALSE;
10     else
11        $\tilde{\Omega} = \text{ADDLAYER}(\tilde{\Omega}, \Omega, \mathcal{C})$ ;
12   }
13   if ( $\tilde{\Omega} \neq \hat{\Omega}$ ) {
14     refinement =  $\text{GETREFFROMCA}(\hat{\Omega}, \tilde{\Omega}, P, L)$ ;
15      $\hat{\Omega} = \text{ADDRIFTOABSMODEL}(\hat{\Omega}, \text{refinement})$ ;
16   }
17    $L = L + 1$ ;
18 }
19 while ( $\text{CHECKSIMPLEPATH}(\Omega, P, L)$ ) {
20   if ( $\text{EXISTCEX}(\Omega, P, L)$ )
21     return FALSE;
22    $L = L + 1$ ;
23 }
24 return TRUE;
}

```

Fig. 4 The PureSAT algorithm

3 Algorithm

In this section we present an algorithm for abstraction refinement that uses SAT as decision procedure on both abstract and concrete models. The input to the algorithm is an open system $\Omega = \langle V, W, I, T \rangle$ whose transition relation T is specified by a circuit graph $\mathcal{C} = (V \cup W, E)$, and a predicate $P(V)$ describing a set of accepting states. We assume that cone-of-influence reduction is applied before invoking the abstraction refinement procedure so that all vertices of \mathcal{C} are in the COI of the property.

3.1 The PureSAT Algorithm

Our algorithm is shown in Fig. 4. Initially, an abstract model $\hat{\Omega}$ is computed by collecting in \hat{V} only the state variables that appear in $P(V)$; hence, $\hat{P} = P$ throughout. In Lines 3–18 the algorithm progressively increases L from its initial value of 0 until either a counterexample of length L is found in the concrete system Ω , or it is concluded that no counterexample exists in the current abstract model $\hat{\Omega}$. If at some point, the abstract model becomes the concrete model, the endgame described in Lines 19–24 is executed.

For each length L , (5) and (6) are checked to see whether the simple path conditions are met. If either one is unsatisfiable, the property holds and the algorithm terminates.

Otherwise, the algorithm checks whether there exists a counterexample of length L (Lines 7–12). The check is incremental: Since every abstract model simulates the concrete

one, a sequence of increasingly refined abstract models is used to establish the existence of a counterexample. If any model $\tilde{\Omega}$ in the sequence has no counterexample of length L , neither does the concrete model Ω . In the first iteration of the inner **while** loop, (3) is checked on the current abstract model $\hat{\Omega}$. If abstract counterexamples exist, function **ADDLAYER** produces a coarse refinement as the next element in the sequence.

The loop terminates as soon as it is known whether a counterexample of length L exists in the concrete model Ω . When no such counterexample exists, if $\tilde{\Omega} \neq \hat{\Omega}$ at Line 13, then the abstract model admits spurious ACEs. Therefore, it is refined by adding to \hat{V} a minimal set of variables from \tilde{V} sufficient to rule out all counterexamples of length L . (Lines 14–15). The resulting model has no counterexamples of length up to L .

Lemma 1 *The abstract model computed at Line 15 of algorithm PURESAT has no counterexample to P of length less than or equal to L .*

Proof Let $\hat{\Omega}_i$ be the abstract model computed at Line 15 when $L = i$. By definition of refinement, $\hat{\Omega}_i$ has no counterexamples of length i . We can then prove by induction that $\hat{\Omega}_L$ has no counterexamples of length i for $i < L$. The base for $L = 0$ is trivially established. For $L > 0$, we can assume that $\hat{\Omega}_{L-i}$ has no counterexample of length less than L . Since $\hat{\Omega}_{L-i}$ is an abstraction of $\hat{\Omega}_L$, a counterexample of length less than L in the latter would have a matching run in the former, which would contradict the inductive hypothesis. \square

3.2 Concretization Test

For each value of L , algorithm PURESAT checks whether the concrete model Ω has a counterexample to P of that length. This test proceeds incrementally, starting from the current abstract model $\hat{\Omega}$, and possibly reaching Ω . At each iteration of the while loop at Lines 7–12 of Fig. 4, a model $\tilde{\Omega}$ is checked. Initially, $\tilde{\Omega} = \hat{\Omega}$; in the worst case, eventually, $\tilde{\Omega} = \Omega$; however, the loop terminates before this condition is met often enough to justify the incremental approach. If a given $\tilde{\Omega} \neq \Omega$ has counterexamples of length L , **ADDLAYER** is invoked to produce the model to be checked at the next iteration.

Procedure **ADDLAYER** adds variables from $V \setminus \hat{V}$ to $\tilde{\Omega}$ with two objectives: To allow the concretization check to complete quickly; and to help produce a better refinement if one is necessary.

When no concrete counterexamples are found, the two goals are met if **ADDLAYER** arrives in few iterations to a small $\tilde{\Omega}$ such that $\text{EXISTCEX}(\tilde{\Omega}, P, L)$ returns false. By producing a short sequence of small models, the time spent in checking for counterexamples in them is kept short. In addition, the sizes of the unsatisfiable cores produced when checking (3) grow with the sizes of the models. A smaller $\tilde{\Omega}$ therefore tends to save time also during refinement minimization. When there are concrete counterexamples, eventually the concrete model Ω must be checked. Hence, only a few iterations of the **while** loop should be taken.

```

set GETREFFROMCA( $\widehat{\Omega}$ ,  $\widetilde{\Omega}$ ,  $P, L$ ) {
1  nsVarSet = GETNEXTSTATEVARSFROMCDG( $\widetilde{\Omega}, P, L$ );
2  RCArray = SORTCANDIDATEARRAY(nsVarSet,  $\widetilde{\Omega}, \widehat{\Omega}$ );
3  sufficient =  $\emptyset$ ;
4  while (sufficient does not kill all length- $L$  counterexamples
5     $\wedge$  RCArray is not empty) {
6    someNsVars = PICKVARS(RCArray, threshold);
7    sufficient = sufficient  $\cup$  someNsVars;
8    RCArray = nsVarSet  $\setminus$  someNsVars;
9  }
10 return REFINEMENTMINIMIZATION( $\widehat{\Omega}$ , sufficient);
}

```

Fig. 5 The refinement algorithm

These requirements are addressed by the following procedure. The variables in $V \setminus \widetilde{V}$ are ranked relative to $\widetilde{\Omega}$ according to the same criteria used in the computation of refinement. If the variables at the lowest level make up less than a set fraction of the variables in $V \setminus \widetilde{V}$, they are all added to the model to obtain the new $\widetilde{\Omega}$. Otherwise, only that fraction is added. Since a variable’s level is the first sorting criterion, all new variables come from the the lowest level; hence, they are guaranteed to have some influence on the behavior of the model.

If the iterations of the **while** loop are too many, the benefit of finding a more compact conflict core does not outweigh the cost of checking the many abstract models for counterexamples. So, if there still are counterexamples after the first few rounds, the algorithm skips to the concrete model.

3.3 Refinement

The goal of the refinement procedure is to find a minimal set of state variables not in $\widehat{\Omega}$ which, when added to the abstract model, can suppress all counterexamples of length L . Care is put into keeping the abstract model small because the success of the termination criterion based on (5) and (6) is affected by the complexity of $\widehat{\Omega}$.

Our algorithm for picking refinement variables is shown in Fig. 5. When GETREFFROMCA is called, $\widehat{\Omega}$ has counterexamples to P of length L , while $\widetilde{\Omega}$ does not because EXISTCEX just returned false. The addition to $\widetilde{\Omega}$ of the state variables in $\widetilde{V} \setminus \widehat{V}$ found in the unsatisfiable core of (3) as next state variables is sufficient to suppress all spurious counterexamples of length L . Variable v_i is used as next state variables in the unsatisfiable core of (3) if w_i^j appears in some clause in that core for some j between 0 and L .

Sufficiency of the refinement can be argued by observing that the candidate variables are all the variables in a set of clauses from which a contradiction can be derived. Refinement with respect to all those variables will add to the SAT problem for the abstract model all the clauses that are necessary to obtain the same contradiction. If state variable v_i does not appear as next state variable in the unsatisfiable core, then

the corresponding term T_i in (1) is irrelevant to the proof of unsatisfiability.

The original “sufficient set” (nsVarSet in Fig. 5) extracted from the unsatisfiable core may be non-minimal; hence, REFINEMENTMINIMIZATION removes redundant state variables from the refinement set. This procedure tentatively removes one variable at the time from the refinement and checks whether EXISTCEX still returns false. Since the order in which candidate variables are considered by the minimization procedure affects the result, SORTCANDIDATEARRAY sorts the variables so that those considered less valuable are removed first.

The number of redundant state variables in nsVarSet may be quite large, causing too many calls to EXISTCEX. The **while** loop of Lines 4–9 is used to heuristically get a smaller “sufficient set” for the refinement minimization: Each time, only a certain number of state variables are picked from RCArray, after which (3) is checked to see if they are already sufficient.

In our method, the order in which state variables are removed in the minimization procedure is based on two criteria: the *level* of each candidate state variable and, as a tie-breaker, its *relative correlation* to the current abstract model. The level of a state variable $v \in \widetilde{V} \setminus \widehat{V}$ w.r.t. $\widehat{\Omega}$ is the least number of state variables on a path in \mathcal{C} connecting v to any state variable in \widehat{V} . The smaller its level, the more important v is deemed. The relative correlation of $v \in \widetilde{V} \setminus \widehat{V}$ equals the ratio of the number of direct predecessors of v in \mathcal{C} that are in $\widehat{\Omega}$ to the total number of nodes in the COI of v . Intuitively, the larger the relative correlation of a state variable, the larger effect it will have when added to or subtracted from the current abstract model. The state variables with the larger levels and smaller relative correlations are considered of less importance, and thus will be tested for deletion earlier. In this way, we can concentrate on the important candidates and keep the refined abstract model small.

3.4 Incremental SAT Solver

An incremental SAT solver exploits the similarities among SAT instances that form a sequence by using the conflict clauses generated from the previous instance to guide the search for a solution to the current instance. We assume that a sequence of SAT instances $(\Delta_0, \dots, \Delta_l)$ is specified by a sequence of pairs of sets of clauses $((A_0, B_0), \dots, (A_l, B_l))$ such that $\Delta_0 = A_0$, $B_0 = \emptyset$, and, for $1 \leq i \leq l$, $\Delta_i = (\Delta_{i-1} \setminus B_i) \cup A_i$. That is, the second and successive instances are obtained by removing some clauses from the instances immediately preceding them, and then adding some other clauses.

If Γ_i is the set of conflict clauses produced while solving Δ_i , $\gamma \in \Gamma_i$ is implied by the conjunction of the clauses in Δ_i . If $B_{i+1} = \emptyset$, it follows that γ is implied also by the conjunction of the clauses in Δ_{i+1} . Adding Γ_i to Δ_{i+1} does not change the answer to the SAT problem, but may help the solver by precluding examination of assignments that are known not to be satisfying. If, on the other hand $B_{i+1} \neq \emptyset$,

$\gamma \in I_i$ may or may not be implied by Δ_{i+1} . However, if none of the clauses that were used in deriving γ is in B_{i+1} , then it is still correct to add γ to Δ_{i+1} . This test is conservative, but inexpensive.

Keeping all valid conflict clauses from previous instances may result in too many clauses with very many literals, which may have a negative impact on search time. In our implementation, we address this problem by modifying the conflict clause deletion strategy: We decrease the interval between successive scans of the database for clauses to be deleted. To avoid wasting time when there are few large conflict clauses generated, we use a threshold. At the end of each interval, we search for clauses to be deleted only if the number of new large conflict clauses exceeds the threshold.

The PURESAT algorithm offers several opportunities to apply an incremental SAT solver. In our implementation, we deal incrementally with two sequences of instances: One is the sequence of checks of (3) (existence of counterexamples), the other is the sequence that interleaves the checks of (5) and (6) (existence of simple paths). In the former, the instances for the same value of L form monotonically increasing subsequences, in which all conflict clauses can be re-utilized. However, when L is incremented, the clauses accumulated for models other than $\hat{\Omega}$ are discarded.

Two more opportunities of applying incremental SAT solving exist in the PURESAT algorithm. In the loop of Lines 5–9 in Fig. 5, and in refinement minimization. However, the average number of iterations of the loop in Fig. 5 is too low to warrant incremental solving, while the SAT instances occurring in refinement minimization are monotonically decreasing; hence, they are not well suited to incremental solving.

4 Related Work

Our refinement algorithm is based on computing and analyzing the unsatisfiable core associated with the proof that there is no concrete counterexample of length L ; hence, it is similar to the conflict analysis method proposed in [CCK⁺02]. However, our approach differs significantly from [CCK⁺02] in the following aspects:

1. The authors of [CCK⁺02] identify a single spurious abstract counterexample (by using BDD-based model checking), together with its *failure index*. (I.e., the time step from which the ACE is no longer concretizable in Ω .) A conflict dependency graph is built from the unsatisfiable formula obtained by constraining the concrete model with the single spurious ACE up to the failure index time step. The refinement set is then computed by analyzing the conflict dependency graph. In our algorithm, however, we do not use a single abstract counterexample to constrain the BMC instance (and consequently we do not compute the failure index). Rather, unconstrained BMC instances (on a sequence of increasingly concrete models) is used for the concretization test; such BMC instances cover all the possible length- L spurious abstract counterexamples.

2. In [CCK⁺02], the *invisible* state variables (those in $V \setminus \hat{V}$) are added to the refinement set if their corresponding literals at the failure index time step appear in the conflict dependency graph. In our algorithm, all the literals (which correspond to either state variables or internal logic gates at different time steps) appearing in the unsatisfiable core are recorded in the SAT solver. However, only those invisible state variables whose *next state variable* literals appear in the unsatisfiable core are added to the refinement set.
3. Our refinement minimization algorithm is also somewhat different from [CCK⁺02]. Both methods remove redundant state variables greedily, but they differ in the order in which variables are considered.

Our algorithm is also related to the one of [MA03]. Algorithm `pureSAT` starts from the abstract model and checks concretization for increasing counterexample lengths. By contrast, the method of [MA03] analyzes the concrete model; for some lengths of counterexamples it builds an abstraction from the failed proof of their existence. Both approaches check all counterexamples of a certain length at once. The main differences are:

1. We use SAT, instead of a BDD-based model checker, for the abstract model. Some abstract models are more suitable to BDD-based model checking; for instance, those that have very long simple paths. Other models have unwieldy BDDs and are more suitable to SAT-based induction proofs. Hence, neither approach dominates the other. As pointed out in Sect. 1, we expect a combination of the two approaches to provide significant advantages.
2. Algorithm `pureSAT` checks concretization incrementally. In many cases, a property can be proved without ever checking the concrete model. An additional advantage of this approach is that unsatisfiable cores come from smaller models. Finally, the abstract model provides guidance in checking the concrete one in the form of conflict clauses of the more abstract models that the incremental SAT solver may use to guide the search of counterexamples in Ω .
3. Our abstraction grows at each refinement, and we use refinement minimization to control its size, whereas the abstraction of [MA03] is computed from scratch each time. Refinement minimization requires repeated BMC runs; these, however, are not on the full concrete model. Re-computing the abstraction from scratch may lead to smaller abstract models. Building it incrementally allows PURESAT to check only for paths of length L instead of length up to L .

In [CGKS02, CCK⁺02, WLJ⁺03], once counterexamples of length L are found on an abstract model, the concrete model is checked immediately to see if a real counterexample of the same length can be found. The method of [MA03] also checks the concrete model for increasing values of L . If the concrete model is very large, this approach may suffer from the state explosion problem. In [BGG02], the authors propose a *layered reconstruction algorithm*, in which, instead

of checking directly the concrete model when abstract counterexamples are found, the state variables in $V \setminus \widehat{V}$ are divided into layers, and as each layer is added in turn, the resulting model is checked for surviving counterexamples of length L . In the PURESAT algorithm, we have adopted a similar layered approach in our incremental concretization test, with the following main differences:

1. While the approach of [BGG02] is completely based on BDDs, our approach is entirely based on SAT. The main consequence of this choice is that we are less concerned about incremental concretization reaching Ω . We do not employ backtracking, and we limit the number of refinements during each concretization test to save time. We can also easily benefit from incremental SAT solving.
2. The computations of the layers differ in the two methods. In particular, in our algorithm we use two criteria: the level of a state variable w.r.t. the current abstract model as the primary criterion, and the relative correlation as the second one.

Application of incremental SAT solving to bounded model checking was proposed in [WKS01]. In [ES03], a new encoding scheme for BMC was presented, which allows efficient use of incrementality. The main observation behind that encoding is that removal of unit clauses can be implemented very efficiently. Our incremental solver does not make use of this observation, and is therefore closer to the one of [WKS01]. Unlike either [WKS01] or [ES03], we use incremental SAT solving for abstraction refinement, rather than plain BMC.

5 Experimental Results

To evaluate the technique of Section 3, we compared four algorithms: Our implementation of the BMC algorithm of [BCCZ99], BMC extended with the checks for simple paths [SSS00] (referred to as SSS), our PURESAT algorithm, and the GRAB algorithm of [WLJ⁺03], which uses both BDDs and SAT. All four algorithms are implemented in VIS-2.0 [B⁺96, VIS]. We used Chaff [MMZ⁺01] as the SAT solver. We added to it the ability to extract unsatisfiable cores and to perform incremental checks. The experiments were run under Linux on an IBM IntelliStation with a 1.7 GHz Intel Pentium 4 CPU and 2 GB of RAM.

The comparison was conducted on 26 models, either from industry or from VIS verification benchmarks [B⁺96, VIS] except for *isp*. This model was created to illustrate the help BMC could get from abstraction. A simplified version of it appears in Fig. 1. Since in the concrete model, the longest simple path is quite long, SSS failed to complete, even though PURESAT finished within one second.

The results are shown in Table 1. The first column is the name of the model; the second column indicates whether each property passes or fails; if a property fails, the number in this column is the length of the counterexample. The third column gives the number of state variables in the cone of influence of the property. The fourth column lists the time of BMC.

A time in parentheses is the time elapsed when the process ran out of memory. In our experiments, the time limit was set to 8 hours. The fifth column is the time of SSS; the sixth column shows the time for GRAB; the seventh column is the number of state variables in the final abstract model. If the time is greater than 8 hours, the number in parentheses in this column is the number of state variables in the abstract model when time ran out. The next two columns are the data for PURESAT. All CPU times are in seconds except when noted.

The algorithm labeled BMC can check inductive invariants. However, no such properties are included in our set of experiments. From the table we can see that, in general, for passing properties, PURESAT is better than both BMC and SSS. For most failing properties BMC is best, while PURESAT is better than GRAB. For the largest models, like IU, whose COI contains 4493 state variables, PURESAT is the only one being able to verify the property. Interestingly, GRAB and PURESAT fail to finish similar numbers of experiments (4 for GRAB and 3 for PURESAT). However, the two sets of failures are disjoint. This is an encouraging sign for the development of a hybrid algorithm that may switch between BDDs and SAT for the analysis of the abstract models.

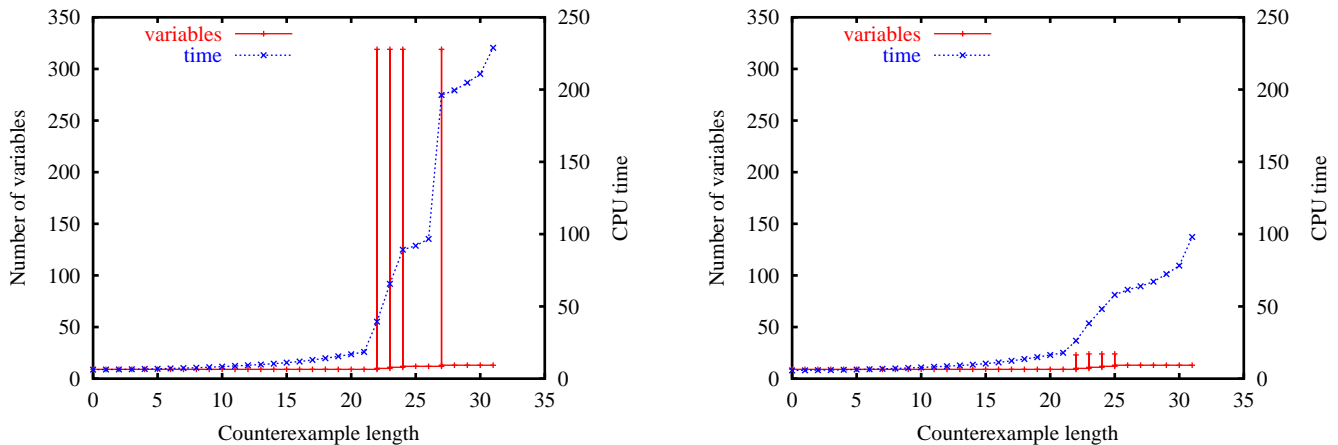
The effects on performance of incremental concretization and incremental SAT solving are shown in Table 2. In this table, IC means incremental concretization, NIC means non-incremental concretization, IS means incremental SAT, NIS means non-incremental SAT. The sizes of the final abstract models are the same for PureSAT(NIC, NIS), PureSAT(IC, NIS), and PureSAT(IC, IS) except for D24-p3 (6, 5, 6), D24-p4 (6, 5, 6), and D24-P5 (6, 8, 8). These small differences are due to the differences in the search order followed by the SAT solver.

From Table 2 one sees that incremental SAT solving almost always decreases CPU time. Incremental concretization helps most of the time, though for some models (e.g., D20-p1) it is not effective because most of the concretization checks end up dealing with Ω . Figure 6 and Table 3 take a closer look at the effects of incremental concretization on model D5-p1. The figure compares two runs, with and without incremental concretization. Each graph shows the numbers of state variables in the models and the CPU times as a function of the counterexample length L . When refinement occurs, both the size of the abstract model and the size of the largest $\widetilde{\Omega}$ examined during the concretization check are shown as the endpoints of a vertical line. For example, the first refinement of the initial abstraction happens for $L = 22$. Without incremental concretization (left graph), all 319 state variables are included in $\widetilde{\Omega}$, whereas the incremental check requires only the analysis of a model with 23 state variables. In both cases, $\widehat{\Omega}$ contains 9 state variables.

Checking smaller models reduces the total CPU time by over 50% even without incremental SAT. The graphs of Fig. 6 clearly show that the time penalty is incurred when refinement takes place. Table 3 examines the nature of this time penalty; it shows that the major effect of incremental concretization is to speed up refinement minimization. This happens because smaller unsatisfiable cores are usually extracted

Table 1 Experimental results. Boldface is used to highlight best CPU times

model	pass/ cex length	latches in COI	BMC time	SSS time	Grab		PureSAT	
					time	final sz.	time	final sz.
lsp-p1	pass	12	NA	>8h	1	3	1	3
D12-p1	16	48	5	25	14	23	22	23
D23-p1	5	85	1	1	20	21	1	25
D2-p1	14	94	6	25	180	48	10	48
D14-p1	14	96	65	83	>8h	(75)	1294	79
D1-p1	9	101	1	5	9	21	6	20
D1-p2	13	101	2	12	51	23	13	23
D1-p3	15	101	3	18	56	25	17	23
I12-p1	370	119	>8h	>8h	2503	16	>8h	(12)
B-p1	pass	124	NA	>8h	173	18	825	18
B-p2	17	124	150	675	93	7	113	7
B-p3	pass	124	NA	>8h	223	43	>8h	(42)
B-p4	pass	124	NA	MO(23708)	393	42	>8h	(42)
D22-p1	10	140	2	10	720	132	10	132
D24-p1	9	147	7	10	1	4	1	4
D24-p2	pass	147	NA	16	3	8	8	8
D24-p3	pass	147	NA	1	20	8	3	6
D24-p4	pass	147	NA	1	43	8	3	6
D24-p5	pass	147	NA	1	3	5	3	8
M0-p1	pass	221	NA	MO(2537)	136	16	1188	13
D5-p1	31	319	58	592	31	18	53	13
D18-p1	23	506	96	795	>8h	(99)	2721	166
D16-p1	8	531	10	29	92	14	20	14
D20-p1	14	562	26	101	>8h	(69)	7445	229
rcu-p1	pass	2453	NA	MO(3115)	195	10	87	10
IU-p2	pass	4493	MO(11331)	>8h	>8h	(6)	1060	14

**Fig. 6** Model sizes and CPU time for D5-p1 without and with incremental concretization (no incremental SAT in either case)

from smaller models. (Practical methods that produce proofs of unsatisfiability do not guarantee proofs that are minimal in terms of either number of clauses or variables appearing in the unsatisfiable core.) Minimization produces refinements of the same sizes for D5-p1 (1 variable in all cases) with and without incremental concretization check. The time advantage comes from examining fewer candidates on smaller models.

In spite of the advantages of dealing with smaller unsatisfiable cores, our experiments indicate that the number of

rounds in the concretization check should be kept small. We obtained the best results when the number of rounds was limited to two.

Though PURESAT, appears to be reasonably robust, there are a few cases in which it performs poorly compared to competing techniques. The profiles of two such runs are shown in Fig. 7.

Model D20-p1 is representative of a group, including also D14-p1 and D18-p1, for which PURESAT does substantially worse than BMC and SSS on failing properties. The left graph

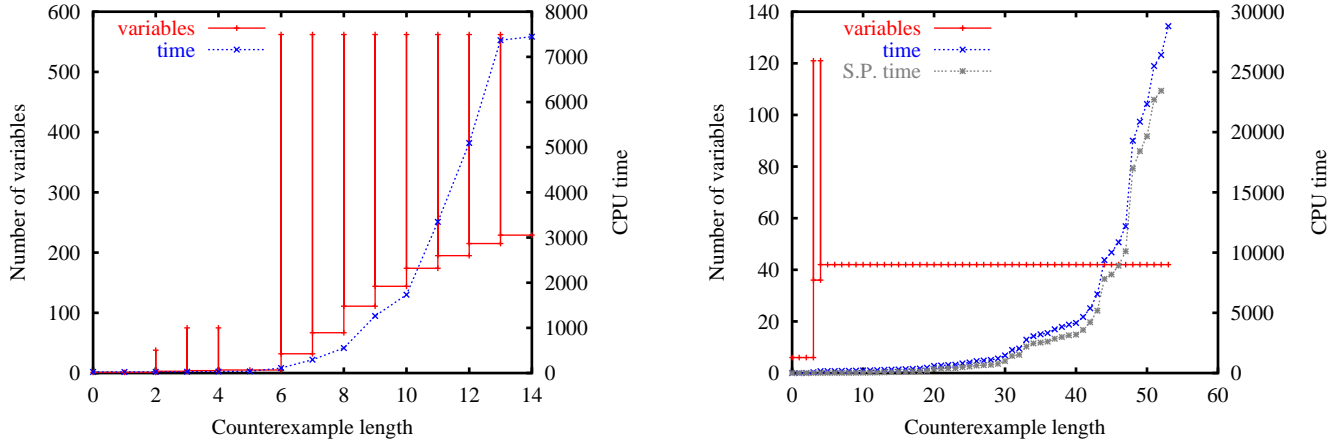


Fig. 7 Two hard problems for PURESAT: D20-p1 (left) and b-p3 (right)

Table 2 Comparison between PureSAT with and without incremental concretization (IC) and with and without incremental SAT (IS). Boldface is used to highlight best CPU times

model	NIC,NIS time	IC,NIS time	IC,IS time
lsp-p1	1	1	1
D12-p1	33	26	22
D23-p1	2	2	1
D2-p1	20	22	10
D14-p1	1393	1345	1294
D1-p1	9	8	6
D1-p2	19	23	13
D1-p3	28	17	13
I12-p1	>8h	>8h	>8h
B-p1	846	857	825
B-p2	316	198	113
B-p3	>8h	>8h	>8h
B-p4	>8h	>8h	>8h
D22-p1	15	18	10
D24-p1	2	2	1
D24-p2	8	9	8
D24-p3	2	5	3
D24-p4	4	2	3
D24-p5	3	3	3
M0-p1	2826	2134	1188
D5-p1	222	98	53
D18-p1	3875	2844	2721
D16-p1	29	24	20
D20-p1	7292	7717	7445
rcu-p1	123	87	87
IU-p2	1683	1568	1060

in Fig. 7 shows that incremental concretization is not effective for D20-p1. As a consequence, refinement minimization is expensive, and accounts for the majority of the CPU time. In this case, the effort put into keeping the abstract model $\widehat{\mathcal{Q}}$ small is wasted: the property is eventually found to fail, and for the values of L up to 31 (the length of the shortest counterexamples), the simple path checks of (5) and (6) are not

Table 3 Comparison of refinement procedures with and without incremental concretization for D5-p1

Ref. Round	Variables in core		Ref. min. time		Concr. time	
	NIC	IC	NIC	IC	NIC	IC
1	17	8	11.24	5.62	3.29	0.51
2	16	8	14.61	7.34	4.64	0.59
3	22	7	19.16	6.51	4.42	0.53
4	50	6	78.69	5.72	7.41	0.61

very expensive even on the concrete model—as witnessed by the CPU time of SSS.

For model b-p3, PURESAT runs out of time, while GRAB proves that the property passes in less than four minutes. The right graph in Fig. 7 shows that PURESAT quickly reaches the stage where no counterexamples are found in the abstract model, but then it is unable to reach a value of L sufficient to prove the property true. In this case, the graph shows also the time (S.P. time) spent in checking (5) and (6), which clearly dominates the total CPU time. It is consistent with this observation that SSS also runs out of time on b-p3. Even though the abstract model produced by PURESAT has one fewer state variable than the one of GRAB, abstraction fails to sufficiently reduce the length of the longest simple paths; hence, a BDD-based approach like the one of GRAB (or [MA03]) is more effective.

Since checking for simple paths is sometimes expensive, one may choose to use only one of (5) and (6). However, in our experiments, both tests contributed to establish termination for passing properties. When the test for (5), which is the most effective of the two, was disabled, the number of experiments successfully completed decreased, and no significant speedups were observed.

6 Conclusions

We have presented an abstraction refinement algorithm for model checking safety properties that uses a SAT solver as

sole decision procedure. We have compared this algorithm to both BMC and to an abstraction refinement algorithm that uses both BDDs and CNF SAT. The new algorithm is competitive and was the only one to complete the largest test case. Our implementation is still preliminary. We are interested in the extension of the techniques of [WLJ⁺03] to the SAT environment. This is not an entirely trivial task, since they are based on the knowledge of the sets of states at various distances along the paths connecting initial states to error states.

By its very nature, the PURESAT algorithm suffers, albeit in attenuated form, from the same problems that afflict the basic procedure used in analyzing the abstract models. Improvements like those proposed in [McM03] may boost PURESAT's performance. Improved extraction of unsatisfiable cores may speed up the abstraction minimization phase, which is currently the most time consuming part of the algorithm. More generally, the integration with a BDD-based approach to the analysis of the abstract model should lead to a more robust and powerful approach to abstraction refinement.

References

- [ABE00] P. A. Abdulla, P. Bjesse, and N. Eén. Symbolic reachability analysis based on SAT-solvers. In *Tools and Algorithms for the Construction of Systems (TACAS)*, pages 411–425, 2000. LNCS 1785.
- [AS85] B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21:181–185, October 1985.
- [B⁺96] R. K. Brayton et al. VIS: A system for verification and synthesis. In T. Henzinger and R. Alur, editors, *Eighth Conference on Computer Aided Verification (CAV'96)*, pages 428–432. Springer-Verlag, Rutgers University, 1996. LNCS 1102.
- [BCCZ99] A. Biere, A. Cimatti, E. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *Fifth International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS'99)*, pages 193–207, Amsterdam, The Netherlands, March 1999. LNCS 1579.
- [BGG02] S. Barner, D. Geist, and A. Gringauze. Symbolic localization reduction with reconstruction layering and backtracking. In E. Brinksma and K. G. Larsen, editors, *Fourteenth Conference on Computer Aided Verification (CAV 2002)*, pages 65–77. Springer-Verlag, July 2002. LNCS 2404.
- [BKA02] J. Baumgartner, A. Kuehlmann, and J. Abraham. Property checking via structural analysis. In E. Brinksma and K. G. Larsen, editors, *Fourteenth Conference on Computer Aided Verification (CAV'02)*, pages 151–165. Springer-Verlag, Berlin, July 2002. LNCS 2404.
- [Bry86] R. E. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, August 1986.
- [CC77] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by constructions or approximation of fixpoints. In *Proceedings of the ACM Symposium on the Principles of Programming Languages*, pages 238–250, 1977.
- [CCK⁺02] P. Chauhan, E. Clarke, J. Kukula, S. Sapra, H. Veith, and D. Wang. Automated abstraction refinement for model checking large state spaces using SAT based conflict analysis. In M. D. Aagaard and J. W. O'Leary, editors, *Formal Methods in Computer Aided Design*, pages 33–51. Springer-Verlag, November 2002. LNCS 2517.
- [CGKS02] E. Clarke, A. Gupta, J. Kukula, and O. Strichman. SAT based abstraction-refinement using ILP and machine learning. In E. Brinksma and K. G. Larsen, editors, *Fourteenth Conference on Computer Aided Verification (CAV 2002)*, pages 265–279. Springer-Verlag, July 2002. LNCS 2404.
- [CGP99] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, Cambridge, MA, 1999.
- [ES03] N. Eén and N. Sörensson. Temporal induction by incremental SAT solving. *Electronic Notes in Theoretical Computer Science*, 89(4), 2003. First International Workshop on Bounded Model Checking. <http://www.elsevier.nl/locate/entcs/>.
- [GN03] E. Goldberg and Y. Novikov. Verification of proofs of unsatisfiability for CNF formulas. In *Design, Automation and Test in Europe (DATE'03)*, pages 886–891, Munich, Germany, March 2003.
- [KS03] D. Kröning and O. Strichman. Efficient computation of recurrence diameters. In *Verification, Model Checking, and Abstract Interpretation*, pages 298–309, New York, NY, January 2003. Springer. LNCS 2575.
- [Kur94] R. P. Kurshan. *Computer-Aided Verification of Coordinating Processes*. Princeton University Press, Princeton, NJ, 1994.
- [MA03] K. L. McMillan and N. Amla. Automatic abstraction without counterexamples. In *International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS'03)*, pages 2–17, Warsaw, Poland, April 2003. LNCS 2619.
- [McM94] K. L. McMillan. *Symbolic Model Checking*. Kluwer Academic Publishers, Boston, MA, 1994.
- [McM02] K. L. McMillan. Applying SAT methods in unbounded symbolic model checking. In E. Brinksma and K. G. Larsen, editors, *Fourteenth Conference on Computer Aided Verification (CAV'02)*, pages 250–264. Springer-Verlag, Berlin, July 2002. LNCS 2404.
- [McM03] K. L. McMillan. Interpolation and SAT-based model checking. In W. A. Hunt, Jr. and F. Somenzi, editors, *Fifteenth Conference on Computer Aided Verification (CAV'03)*, pages 1–13. Springer-Verlag, Berlin, July 2003. LNCS 2725.
- [Mil71] R. Milner. An algebraic definition of simulation between programs. *Proc. 2nd Int. Joint Conf. on Artificial Intelligence*, pages 481–489, 1971.
- [MMZ⁺01] M. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *Proceedings of the Design Automation Conference*, pages 530–535, Las Vegas, NV, June 2001.
- [SSS00] M. Sheeran, S. Singh, and G. Stålmarck. Checking safety properties using induction and a SAT-solver. In W. A. Hunt, Jr. and S. D. Johnson, editors, *Formal Methods in Computer Aided Design*, pages 108–125. Springer-Verlag, November 2000. LNCS 1954.
- [VIS] URL: <http://vlsi.colorado.edu/~vlsi>.
- [WBCG00] P. Williams, A. Biere, E. M. Clarke, and A. Gupta. Combining decision diagrams and SAT procedures for

- efficient symbolic model checking. In E. A. Emerson and A. P. Sistla, editors, *Twelfth Conference on Computer Aided Verification (CAV'00)*, pages 124–138. Springer-Verlag, Berlin, July 2000. LNCS 1855.
- [WHL⁺01] D. Wang, P.-H. Ho, J. Long, J. Kukula, Y. Zhu, T. Ma, and R. Damiano. Formal property verification by abstraction refinement with formal, simulation and hybrid engines. In *Proceedings of the Design Automation Conference*, pages 35–40, Las Vegas, NV, June 2001.
- [WKS01] J. Whittemore, J. Kim, and K. Sakallah. SATIRE: A new incremental satisfiability engine. In *Proceedings of the Design Automation Conference*, pages 542–545, Las Vegas, NV, June 2001.
- [WLJ⁺03] C. Wang, B. Li, H. Jin, G. D. Hachtel, and F. Somenzi. Improving Ariadne's bundle by following multiple threads in abstraction refinement. In *Proceedings of the International Conference on Computer-Aided Design*, pages 408–415, November 2003.
- [ZM03] L. Zhang and S. Malik. Validating SAT solvers using an independent resolution-based checker: Practical implementations and other applications. In *Design, Automation and Test in Europe (DATE'03)*, pages 880–885, Munich, Germany, March 2003.