

Efficient Coverability Analysis by Proof Minimization^{*}

Alexander Kaiser¹, Daniel Kroening¹, and Thomas Wahl²

¹ University of Oxford, United Kingdom

² Northeastern University, Boston, United States

Abstract. We consider multi-threaded programs with an unbounded number of threads executing a finite-state, non-recursive procedure. Safety properties of such programs can be checked via reduction to the *coverability problem* for *well-structured transition systems* (WSTS). In this paper, we present a novel, sound and complete yet empirically much improved solution to this problem. The key idea to achieve a compact search structure is to track *uncoverability* only for *minimal* uncoverable elements, even if these elements are not part of the original coverability query. To this end, our algorithm examines elements in the downward closure of elements backward-reachable from the initial queries. A downside is that the algorithm may unnecessarily explore elements that turn out coverable and thus fail to contribute to the proof minimization. We counter this effect using a forward search engine that simultaneously generates (a subset of all) coverable elements, e.g., a generalized Karp-Miller procedure. We demonstrate in extensive experiments on C programs that our approach targeting minimal uncoverability proofs outperforms existing techniques by orders of magnitude.

1 Introduction

In anticipation of the prominent role concurrency is predicted to play in future software, popular systems languages like C and Java readily embrace support for multiple threads of execution. Communication among threads is naturally enabled via shared variables, mutexes, but also via non-blocking sleep/wake-up mechanisms. The correct use of these communication mechanisms is largely up to the user. The inevitable frustration caused by attempts to find and reproduce concurrency bugs through conventional program testing strongly encourages the use of automated formal techniques.

In this paper, we consider finite-state, non-recursive procedures executed by an *unspecified* number of threads. This scenario is highly relevant in practice. For example, the number of processes concurrently requesting I/O services in an operating system environment cannot be determined a priori. For settings like this, we are interested in detecting, or proving the absence of, assertion failures, mutual-exclusion violations, etc.

Despite the arbitrary number of threads, problems of this kind have long been known to be decidable [2], for instance by reduction to the *coverability problem* for the rich class of *well-structured transition systems* (WSTS) [23,16]. “Coverability” of an erroneous configuration of threads (e.g., violating mutual exclusion) is tantamount to the existence of a reachable program state exhibiting such an error.

^{*} This research is supported by the EPSRC project EP/G026254/1 and ERC project 280053.

While decidable, checking coverability for WSTS incurs a high computational cost. For example, for the subclass of *vector addition systems* the problem was shown to be EXPSPACE-complete [6]. Extensions such as *transfer transitions*, which allow several threads to change their local state simultaneously and are essential to model broadcast primitives and predicate abstractions of broadcast-free programs [10,9], render the problem Ackermann-hard [33]. The significance of the coverability problem both as a theoretical challenge, as well as in practical concurrent program verification, has led to a flurry of related activity in recent years [29,2,21,20,3,25]. The most general solution to the coverability problem was presented in a paper by Abdulla et al. [2], which backward-explores states starting from the target state.

In this paper, we introduce a new, sound and complete solution to the coverability problem in WSTS. In contrast to existing techniques, our method relies on sequences of many inexpensive uncoverability proofs that build upon one another. We compute such proofs by searching the downward-closure of states encountered during backward exploration from the target state for smallest uncoverable members. Elements encountered during that search that are not currently known to be coverable give rise to “uncoverability candidates”. If a candidate proves uncoverable, so are all elements in its upward closure, which in the end contributes to the decision for the target state. Otherwise, the coverable candidate is retained to prevent it from being expanded again.

The downside of such a “speculative exploration” is that *coverable* exploration candidates mean wasted effort. This effort can, however, be largely reduced using a simultaneously running forward search engine that labels states as coverable and thus prevents them from being explored in the (futile) hope of finding an uncoverability proof. The key is that such a forward engine acts only as a “catalyst”: it affects the speed of the overall algorithm, not its result. Thus, we can use incomplete procedures such as generalizations of the (forward-directed) Karp-Miller algorithm [26,11,14].

To summarize, this work makes the following contributions:

- We present a novel approach to coverability checking in WSTS that combines forward propagation of underapproximations with backward propagation of overapproximations.
- We provide an implementation (publicly available; see Section 5) that accepts Petri nets with transfer arcs, and an extension to verifying C programs with unbounded thread counts in a predicate abstraction-based CEGAR loop [4,7]. Our algorithm outperforms the best known coverability approach by orders of magnitude, enabling the analysis of programs which are out of scope of the previous technology. The experiments also reveal that our approach is able to guide the search far more effectively than existing structural invariant heuristics [13,8].

These improvements are possible thanks to the compactness of the uncoverability proofs generated during exploration. On our C benchmarks, we observe reductions of up to 95% in the proof size.

2 System Model and Problem Definition

Our algorithms operate on *well-structured transition systems* (WSTS) [16]. A WSTS is a transition system equipped with a well-quasi-ordering \succeq on its states that satisfies the following monotonicity property: for all states u, u', r , if u' is a successor of u and $r \succeq u$, then there exists a successor r' of r such that $r' \succeq u'$. In other words, \succeq is a simulation relation for the transition system. A state q is *coverable* if there exists a reachable state v such that $v \succeq q$; the definition of “reachable (with respect to a set of initial states)” is standard.

Let now (M, \succeq) be a WSTS. We denote by Cover the *coverability set*, consisting of all states covered by some reachable state. The *coverability problem* is: given a state $q \in V$, determine whether q is coverable.

Thread Transition Systems. We give an example of a class of WSTS called *thread transition system* (TTS) that are motivated by the task of verifying multi-threaded asynchronous software. A TTS is a machine model that gives rise to transition systems equal in expressiveness to Petri nets [25,28]. We use TTS in examples throughout this paper.

Let S and L be finite sets of *shared* and *local* states, respectively. The elements of $T = S \times L$ are called *thread states*. Formally, a *thread transition system* (TTS) is a pair (T, Δ) , where $\Delta \subseteq T \times T$ models thread transitions. Let $V = \cup_{n=0}^{\infty} (S \times L^n)$. The elements of V are called *states*; we write them in the form $(s \mid l_1, \dots, l_n)$. A TTS gives rise to a transition system $M = (V, \rightarrow)$ with

$$(s \mid l_1, \dots, l_n) \rightarrow (s' \mid l'_1, \dots, l'_n)$$

exactly if, for some $i \in \{1, \dots, n\}$, $(s, l_i, s', l'_i) \in \Delta$ and for all $j \neq i$, $l_j = l'_j$. That is, transitions may affect the shared state, and the local state of exactly one thread in local state l .

Given sets $I_s \subseteq S$ and $I_l \subseteq L$ of initial shared and local states, respectively, we define the set of initial states to be $I = I_s \times (\cup_{n=0}^{\infty} I_l^n)$. An *execution* of transition system M is a finite or infinite sequence of states in V whose adjacent states are related by \rightarrow ; the first state must be initial. A state is *reachable* if it appears in some execution.

To show that M is a WSTS, let the *covers* relation \succeq over V be defined as follows:

$$(s \mid l_1, \dots, l_n) \succeq (s' \mid l'_1, \dots, l'_n)$$

whenever $s = s'$ and $\langle l_1, \dots, l_n \rangle \supseteq \langle l'_1, \dots, l'_n \rangle$, where $\langle x \rangle$ denotes the *multiset* with the elements from x . Let further $v \succ v'$ whenever $v \succeq v'$ and $v \neq v'$. If $0 \in S$ and $0, 1, 2 \in L$, then for example $(0 \mid 0, 2, 0, 1) \succeq (0 \mid 2, 0)$, but $(0 \mid 0, 2, 0, 1) \not\succeq (0 \mid 0, 2, 0, 0)$. Relation \succeq is neither symmetric nor anti-symmetric: states that cover each other are identical *up to permutations of the threads* and thus form a classical thread symmetry equivalence class. Relation \succeq is thus a quasi-order, and in fact a *well-quasi-order* (wqo) on V : any infinite sequence v_1, v_2, \dots of elements from V contains an increasing pair $v_i \preceq v_j$ with $i < j$. It is easy to see that (M, \succeq) fulfills the definition of a WSTS.




Fig. 1. Standard and minimal uncoverability proof for target q . Arrows \rightarrow visualize covering predecessor relations, subscripted by the inducing thread transition; note that $p \rightarrow w$ implies $p \in \text{CPre}(w)$. Arrows $p \succ w$ indicate that there exists some $v \in \text{CPre}(w)$ such that $v \succ p$.

3 Compact Backward Reachable Sets

We introduce some basic definitions and sketch the idea underlying our approach. A set $P \subseteq V$ of states is *upward-closed* if, for any $v \in P$ and any $v', v' \succeq v$ implies $v' \in P$. We write $\uparrow P$ for the *upward closure* of P , i.e., the least upward-closed set that contains P , and $\min P$ for the set of minimal elements of P , i.e., the least subset M of P such that $\uparrow P = \uparrow M$. Every upward-closed set is representable by its minimal elements, of which only a finite number exists due to the wqo properties of \preceq . The term and symbol *downward-closed* and $\downarrow P$ are defined analogously.

Let the *covering predecessors* of a state $v \in V$, denoted by $\text{CPre}(v)$, be all the minimal states that have a successor covering v :

$$\text{CPre}(v) = \min\{p \in V \mid \exists v' \in V : p \succ v' \wedge v' \succeq v\}.$$

Note that for TTS a state $(s \mid \dots) \in \text{CPre}(v)$ involves an additional thread if no thread in v is responsible for the transition to shared state s (we will see an example later on).

Backward Reachability. Algorithm 1 shows a refined version of the classical backward search for WSTS [2,1]. Input is a target state $q \in V$. The algorithm maintains a set $U \subseteq V$ with *vertices* that are labeled and identified with encountered states, and a *work set* $W \subseteq U$ of yet unprocessed vertices.

The algorithm performs an iterative search over covering predecessors starting from q . It terminates either by finding an execution leading to a state that covers q , or when no minimal and unprocessed vertex remains (this eventually happens since \succeq is a wqo), thus proving the uncoverability of the target state.

Algorithm 1. $\text{Bc}(q \in V)$

```

1:  $W := \{q\}; U := \{q\}$ 
2: while  $\exists w \in W : w \in \min(U)$  do
3:    $W := W \setminus \{w\}$ 
4:   for all  $p \in \text{CPre}(w) : p \notin \uparrow U$  do
5:     if  $p \in I$  then
6:       return “ $q \in \text{Cover}$ ”
7:      $W := W \cup \{p\}; U := U \cup \{p\}$ 
8: return “ $q \notin \text{Cover}$ ”
    
```

Minimal Uncoverability Proofs. Let us assume for the rest of this section that the target q is uncoverable. In this case Algorithm 1 computes a representation (in terms of minimal states in U) of all states that have an emanating execution leading to a state that covers q . This set, which we denote by Brs , is upward-closed due to monotonicity.

Instead of computing this set precisely we can, however, also prove the target state uncoverable by any overapproximation Brs^\sharp of this set that still enjoys disjointness from the initial states. The crux is, given that upward-closed sets are represented by their minimal elements, overapproximating these sets by adding smaller (“ \preceq ”) elements leads to fewer and smaller minimal elements, hence to a **more succinct representation**. The following lemma reveals how to exploit this property and settle the uncoverability more efficiently.

Lemma 1. *Let v and v' be two states such that $v' \preceq v$. Then for all $m \in \text{CPre}(v)$ there exists $m' \in \text{CPre}(v')$ such that $m' \preceq m$.*

The proof of Lemma 1 is straightforward. Applied to Algorithm 1 this property generalizes to paths of arbitrary length through the search: the smaller the covering predecessors we examine in Line 4, the shorter the paths we need to explore.

Definition 2. *An **uncoverability proof** for an element q is an upward-closed set of states Brs^\sharp such that $q \in \text{Brs}^\sharp$, $\text{Brs}^\sharp \supseteq \text{CPre}(\text{Brs}^\sharp)$, and $\text{Brs}^\sharp \cap I = \emptyset$. An uncoverability proof Brs^\sharp for q is **minimal** if $\min \text{Brs}^\sharp \subseteq \min(V \setminus \text{Cover})$, and every upward-closed subset $X \subset \text{Brs}^\sharp$ is not an uncoverability proof for q .*

Minimal uncoverability proofs thus consist solely of smallest uncoverable states and cannot be reduced by removing the upward-closure of some of its minimal states. Note that multiple minimal uncoverability proofs may exist.

Bearing Lemma 1 in mind, we observe that minimal uncoverability proofs are an interesting means for proving the uncoverability of target q , as they minimize the maximum length of paths Algorithm 1 needs to traverse.

Example. To illustrate this idea, let us consider the TTS with shared and local states $0, \dots, 3$ and thread transitions $t_1 = (2, 2, 3, 0)$, $t_2 = (0, 2, 2, 0)$, $t_3 = (1, 2, 0, 0)$, $t_4 = (1, 1, 1, 2)$, and $t_5 = (0, 0, 1, 1)$; the initial shared and local state sets are $I_s = I_l = \{0\}$. Assume we wish to check whether shared state 3 can be reached, i.e., the target is $(3 \mid)$. Figure 1 (left) depicts the minimal states of the corresponding set Brs computed by Algorithm 1. If we search, however, the downward-closure of encountered states for smallest uncoverable members, we obtain the minimal uncoverability proof shown on the right: the covering predecessors $(2 \mid 2)$ and $(0 \mid 0, 1, 2)$ give rise to candidate $(2 \mid)$ and $(0 \mid 1)$, respectively. Comparing both uncoverability proofs, we observe reductions in various dimensions: the number of minimal states drops from 9 to 7, the longest traversed path from 7 to 6, and the maximum thread count from 3 to 2. \square

In Section 5 we present experimental evidence that show the potential for compressing the proof size along these dimensions in practice: for our concurrent C program benchmarks we observed **reductions by 95%, 67%, and 50%**, respectively. This potential is the key for the efficiency of our approach.

4 Minimal Uncoverability Proof Algorithm

In this section, we develop our approach to compute minimal uncoverability proofs. An obstacle is the determination of “helpful” candidates. We begin by illustrating it on the TTS from the previous example; we omit non-minimal states for sake of brevity.

Example. Again, we start from target $(3 \mid)$. However, before exploring covering predecessors, we check whether a helpful candidate for a smaller state exists. Since this is not the case (no smaller state exists), we proceed as usual and obtain predecessor $(2 \mid 2)$ which gives rise to candidate $(2 \mid)$. If we find a path showing $(2 \mid)$ coverable, we will withdraw the candidate and proceed with the former state as usual. From $(2 \mid)$ we encounter predecessor $(0 \mid 2)$; although $(0 \mid 2)$ strictly covers $(0 \mid)$, we do not create a corresponding candidate as it is initial. However, for its predecessor $(1 \mid 2, 2)$ in turn we do so, and create candidate $(1 \mid)$. Further exploring this state we obtain path $(0 \mid 0) \rightarrow (1 \mid)$, proving the candidate’s coverability. We withdraw the candidate and mark the downward-closure of all states along the execution as coverable, so that these elements are not expanded again. With path $(0 \mid 0) \rightarrow (1 \mid 1) \rightarrow (1 \mid 2)$ the next candidate proof attempt also fails. From the original predecessor $(1 \mid 2, 2)$ we arrive at $(1 \mid 1, 2)$, of which we can rule out the existence of smaller uncoverable states from the collected coverability results; the same holds for the next predecessor $(1 \mid 1, 1)$. We finally arrive at predecessor $(0 \mid 0, 1)$ and create the candidate $(0 \mid 1)$. Since no new (w.r.t. \uparrow) predecessor exists, we terminate with the tree shown in Figure 1 on the right. \square

4.1 Backward-Constructed Minimal Proofs

In addition to the data structures used by Algorithm 1, namely a set $U \subseteq V$ with vertices that are labeled and identified with encountered states, and a work set $W \subseteq U$ of unprocessed vertices, our algorithm maintains

- i) a set E storing (directed) *edges* between vertices, $E \subseteq U \times U$;
- ii) a mapping ζ *associating* each vertex with a unique vertex, $\zeta: U \rightarrow U$;
- iii) a downward-closed set D storing collected *coverability results*, $D \subseteq V$.

As already indicated in Figure 1, we write $u \rightarrow r$ for $(u, r) \in E$, and \rightarrow^* for the reflexive transitive closure of \rightarrow . We call a vertex $u \in U$ *candidate vertex* if $\zeta(u) = u$, and *predecessor vertex* otherwise. A *path* of (U, E) is a finite sequence of vertices from U whose adjacent vertices are related by \rightarrow ; the *last* state must be a candidate vertex. The mapping ζ (extended to sets X by $\zeta(X) = \{\zeta(x) \mid x \in X\}$) clusters the vertices into $|\zeta(U)|$ *partitions*, one per candidate vertex (vertices that are associated with that candidate vertex). The set D stores states that were shown to be coverable.

The algorithm takes a target q as input and ensures at all times that restricting the partitioned graph (U, E, ζ) to any equivalence class of vertices with the same associated candidate vertex, say u , forms a tree with u as root, and all other vertices being predecessor vertices. Each tree represents an attempt to prove the corresponding candidate (as done by Algorithm 1 for input u). Edges and the mapping ζ enable the withdrawal of unhelpful candidates in a way that preserves parts of their partition that are shared with remaining candidates.



Fig. 2. Effect of routine Backtrack in the presence of candidate vertices s, r and t , each with a single primed predecessor vertex in their partition; the partition of candidate vertex r we wish to remove is highlighted, and the single conflicting edge is marked with “o” (left). After the call, the partition of r is removed and its former predecessor vertex r' is associated with s (right).

The algorithm consists of three routines: Enlarge creates a new candidate vertex, Backtrack removes partitions of candidates, and Mcov is the main routine.

Enlargement Routine. The Enlarge routine takes a candidate u we wish to add as input. If u is a new vertex ($u \notin U$), it is inserted in the work and vertex set. In all cases, the graph is repartitioned by adjusting ζ and associating every vertex in the set

$$\Lambda(u) = \{r \in U \mid u = r \vee (r \rightarrow^* u \wedge \zeta(r) = \zeta(u))\}$$

with u . This repartitioning (observe $u \in \Lambda(u)$) ensures that $r \in \Lambda(u)$ now entails $\zeta(r) = u$. The graph thus contains the new candidate vertex u , with a partition in the shape of a tree.

Backtracking Routine. The purpose of the Backtrack routine, shown in Algorithm 2, is to remove unhelpful candidate vertices $P \subseteq \zeta(U)$ and their partitions. An obstacle is that paths $u \rightarrow^* r \notin P$ to remaining candidate vertices may have segments in partitions that will be removed (paths can traverse multiple partitions). To ensure soundness, we need to preserve them.

Algorithm 2. Backtrack($P \subseteq \zeta(U)$)

- 1: **while** $\exists (r, s) \in E : (r, s)$ is P -confl. **do**
- 2: **for all** $t \in \Lambda(r)$ **do**
- 3: $\zeta(t) := \zeta(s)$
- 4: **for all** $r \in U : \zeta(r) \in P$ **do**
- 5: $W := W \setminus \{r\}; U := U \setminus \{r\}$
- 6: **for all** $(t, r) \in E$ **do**
- 7: $E := E \setminus \{(t, r)\}$

Definition 3. Consider a set P of candidate vertices. An edge $(r, s) \in E$ is called *P-conflicting* if $\zeta(r) \in P$ and $\zeta(s) \notin P$.

Hence, P -conflicting edges induce segments of the above kind. To preserve them, we exhaustively resolve conflicts in a first step (Lines 1–3): for a conflicting edge, say $r \rightarrow s$, we do this by reassociating vertices in $\Lambda(r)$ to $\zeta(s)$.

Once all conflicts are resolved and thus $r \rightarrow s$ and $\zeta(r) \in P$ entails $\zeta(s) \in P$, remaining vertices and edges of partitions in P are removed in Lines 4–7. Figure 2 sketches both steps.

Main Routine. We introduce some terminology:

Algorithm 3. Minimal Uncoverability Proof Algorithm: $\text{Mcov}(q \in V)$

```

1:  $W := \{q\}; U := \{q\}; D := I; E := \emptyset; \zeta : q \mapsto q$ 
2: select  $n \in \min \mathcal{C}(q)$ ; Enlarge( $n$ ) // create candidate vertex
3: while  $\exists w \in W : w \in \min(U)$  do
4:    $W := W \setminus \{w\}$ 
5:   for all  $p \in \text{CPre}(w)$ :  $p$  is  $\zeta(w)$ -minimal do
6:     if  $p \notin D$  then
7:        $E := E \cup \{(p, w)\}$ 
8:       if  $p \notin U$  then
9:          $W := W \cup \{p\}; U := U \cup \{p\}; \zeta(p) := w$  // add covering predecessor
10:        select  $n \in \min \mathcal{C}(p)$ ; Enlarge( $n$ ) // create candidate vertex
11:        else if  $q \notin \downarrow p$  then
12:           $D := D \cup \downarrow p$  // mark coverable states
13:          Backtrack( $\zeta(\downarrow p)$ ) // call backtrack routine
14:          while  $\exists u \in \min(U) \cap \uparrow P$  do
15:            select  $n \in \min \mathcal{C}(u)$ ; Enlarge( $n$ )
16:            break // skip forward to next iteration of while
17:          else
18:            return " $q \in \text{Cover}$ "
19: return " $q \notin \text{Cover}$ "

```

Definition 4. Let $v \in V$, and $u \in \zeta(U)$. State v is *u -minimal* if $v \not\preceq u$ and for all $s, s' \in U$ such that $s \rightarrow s'$ and $\zeta(s') = u$, we have $v \not\preceq s$.

That is, state v is u -minimal if it covers neither the candidate vertex u nor any predecessor vertex in u 's partition (observe that a predecessor vertex may yet belong to a partition other than $\zeta(u)$).

Definition 5. Let $P \subseteq V$. P is *lower successor-closed* if, for any $p \in P$ and any v , $(p \rightarrow v \vee p \succeq v)$ entails $v \in P$.

That is, a lower successor-closed set is both "successor-closed" (where successors are formed according to \rightarrow) and downward-closed. We write $\downarrow v$ for the least lower successor-closed set containing v . This set is obtained by closing $\{v\}$ under \rightarrow successors and downward until fixpoint. The point of this definition is that, if v is coverable, so is every vertex in $\downarrow v$: coverability itself is closed under \rightarrow successors and downward.

Algorithm 3 shows the main routine, Mcov , of our approach. The algorithm works as follows. Initially W and U contain one candidate vertex (target q), D is the set of initial states, the set E of edges is empty, and ζ associates q to itself (Line 1). If target q gives rise to a candidate we create a minimal candidate vertex (Line 2). The set of potential candidates $\mathcal{C}(p) \subseteq V$ is given by

$$\mathcal{C}(p) = \{v \in V \mid v \prec p \text{ and } v \notin D\}.$$

The set contains all the states that are strictly covered by p but not yet marked coverable. If $p = (0 \mid 0, 0, 1)$, and $D = \{(0 \mid), (0 \mid 0), (0 \mid 1)\}$, then for example $\mathcal{C}(p) = \{(0 \mid 0, 1), (0 \mid 0, 0)\}$. We tacitly assume that Line 2 has no side-effect if $\mathcal{C}(p) = \emptyset$.

The algorithm now picks and removes a minimal and unprocessed vertex w from the work set, or returns “ $q \notin \text{Cover}$ ” (Line 19) if no such vertex remains. In the former case, the **for** loop in Line 5 steps through all covering predecessors p of w that are $\zeta(w)$ -minimal and processes them as follows:

Lines 6–10 If p is not currently known to be coverable, then the graph is expanded. If p is a new vertex ($p \notin U$, Line 8), then we ensure that p will be processed when it turns minimal among the vertices by adding it as predecessor vertex to w 's partition. Finally, we call the Enlarge routine to create new minimal candidate vertices.

Lines 11–16 If p is found to be coverable but not q , we add $\downarrow p$ (which is coverable as well) to D and invoke the Backtrack routine to remove partitions of coverable candidate vertices. Since this may remove candidate vertices of remaining predecessor vertices, we have to ensure that their downward-closure is further searched for minimal, yet helpful candidates. We therefore create new minimal candidate vertices (Lines 14–15). Again, we tacitly assume that Line 15 has no side-effect if $\mathcal{C}(p) = \emptyset$. Then, the **break** instruction skips forward to the next iteration of the **while** loop. As a consequence of backtracking, unprocessed vertices that were previously *not* minimal may now be.

Lines 17–18 Otherwise we return “ $q \in \text{Cover}$ ”, since the coverability of target q has been settled (in the affirmative).

Example. We continue with the example from the beginning of this section. In this case routine Enlarge is called four times: predecessor vertices $(2 \mid 2)$, $(1 \mid 2, 2)$, and $(0 \mid 0, 1)$ give rise to candidates $(2 \mid)$, $(1 \mid)$ (and after its removal to $(1 \mid 2)$), and $(0 \mid 1)$, respectively. Routine Backtrack is called once after candidate vertices $(1 \mid 2)$ and $(1 \mid 2, 2)$ turn out unhelpful. The mapping ζ shown in Figure 1 on the right has three partitions, one for each of the candidate vertices $(3 \mid)$, $(2 \mid)$, and $(0 \mid 1)$. The collected coverability results are $D = \downarrow\{(0 \mid 0), (1 \mid 1), (1 \mid 2)\}$, and the mapping ζ is: $\zeta(u) = (2 \mid)$ if $u \in \{(0 \mid 2), (1 \mid 2, 2), (1 \mid 1, 2), (1 \mid 1, 1)\}$, and $\zeta(u) = u$ if $u \in \{(3 \mid), (2 \mid), (0 \mid 1)\}$. \square

Due to the finiteness of downward closures (we create a finite number of candidate vertices) the algorithm eventually terminates. Completeness follows from that of Algorithm 1, and the fact that we only remove conflicting edges during backtracking. When Mcov terminates for an uncoverable target q , the remaining minimal nodes represent an uncoverability proof for q : $\text{Brs}^\# = \uparrow U$ (cmp. Definition 2).

In its current form Algorithm 3 computes uncoverability proofs with the property $\min \text{Brs}^\# \subseteq \min(V \setminus \text{Cover})$, but not necessarily minimal ones. This is attributed to two factors. First, if a covering predecessor gives rise to a candidate and we later remove this predecessor, then a created uncoverability candidate may turn irrelevant for the coverability of target q . Second, when we add a candidate vertex that is incomparable to existing candidate vertices, this may still turn some of the latter irrelevant as well. In order to obtain truly minimal uncoverability proofs, we remove candidate vertices that are no longer needed during calls to Backtrack, and after every call to Enlarge.

4.2 Balancing the Search via Supplementary Coverability Results

If candidates are chosen unwisely, the search may incur extra work to identify and eliminate the coverable elements. To reduce this overhead, we have to prevent unhelpful candidates from being created. In its current form, Algorithm 3 does so by incorporating collected coverability results when it creates a new candidate. These coverability results may also, however, come from any external source, which we call a *coverability oracle*. A coverability oracle a) needs to report states that are provably coverable and should thus reasonably search in a *forward direction*; b) is not required to find *all* coverable states: creating some unhelpful candidates does not harm the search. This flexibility allows us to use any underapproximating forward-directed search: a standard or random reachability analysis works just as well as generalizations of the Karp-Miller procedure to broadcast synchronization [11], which are known not to guarantee termination for WSTS.

We finally remark: since detecting coverable elements is one of the main goals of Algorithm 3, the coverability results reported by the coverability oracle directly benefit the algorithm itself. The coverability oracle and Algorithm 3 run in parallel and synchronize via the set D : the coverability oracle populates this set while maintaining $D \subseteq \text{Cover}$. Receiving such updates, Algorithm 3 terminates if $q \in D$, or otherwise invokes the Backtrack routine on now known-to-be-coverable candidate vertices in regular intervals to restore the invariant $D \cap U = \emptyset$.

5 Experimental Evaluation

In this section, we evaluate our algorithms on 21 concurrent C programs. The programs feature a diverse set of communication primitives, such as shared variables, mutexes, condition variables and broadcasts. For each benchmark, we consider verification of a safety property, specified via an assertion. The C programs, ranging from 40 to 1000 lines of code, are:

- 1–4** broadcast-based code from FreeBSD, NetBSD and Solaris that is related to RDMA ZFS file system support and interface/system monitoring;
- 5–9** programs using several basic language features and the pthread library;
- 10–12** programs using multiple locks to control access to a shared resource;
- 13,14** blocking and non-blocking pseudo-random number generators [31,10];
- 15** a program used in [17] to illustrate thread-modular model checking [24];
- 16,17** lock-based and lock-free stack described in [31], supporting concurrent pushes and pops (adapted from an IBM implementation) [10];
- 18,19** a Linux driver skeleton and a Mozilla vulnerability fix [27,24];
- 20,21** algorithms to establish mutual exclusion [24].

We implemented our Mcov routine (Algorithm 3) for TTSs and transfer Petri nets in our tool BREACH, equipped with a generalization of the Karp-Miller procedure (GKM) as coverability oracle; our tool (we used v1.0) and all benchmarks are available online at www.cprover.org/bfc. The oracle reports coverability results to a data pool our Mcov routine taps into at regular intervals; both run in parallel. In order to measure the impact

Table 1. Comparison of classical coverability approaches to our MCOV algorithm; buggy benchmarks in **bold**, run times in seconds, or **TO (MO)** in case the time (memory) limit is hit

C Programs id/Name	Final TTS		Classical approaches			Our new approach			
	[T]	\Delta	GKM Time	BC Iter.	Time	MCOV		MCOV/GKM	
						Iter.	Time	Iter.	Time
1/BSD-ABDD	82	288	MO	23476	19.1	328	0.1	184	0.0
2/BSD-RDMA-ADDR	101	304	1.6	12479	7.6	295	0.1	146	0.0
3/NETBSD-SYSMON-PWR	291	704	MO	–	TO	124	0.1	126	0.0
4/SOLARIS-SPACE-MAP	539	992	MO	10348	5.8	3412	2.2	2834	1.0
5/BS-LOOP	11616	20485	0.1	1483	1.5	1049	1.1	–	0.1
6/COND	280	1045	0.0	809	0.2	4660	88.4	–	0.0
7/FUNCTION-POINTER	9216	746770	MO	–	TO	–	TO	23139	592.0
8/S-LOOP	516	2813	0.0	3567	1.5	1567	1.4	–	0.5
9/PTHREAD	17920	135300	MO	–	TO	70841	1521.0	51265	189.7
10/DOUBLE-LOCK1	34880	233025	MO	–	TO	–	MO	90488	1146.5
11/DOUBLE-LOCK2	17216	114752	MO	–	TO	–	MO	46012	285.9
12/DOUBLE-LOCK3	3264	19250	MO	–	TO	24161	75.8	9514	14.5
13/PRNG (NON-BL.)	142	954	MO	191	0.0	4791	6.9	64	0.0
14/PRNG	788	5650	MO	–	TO	–	TO	9168	33.9
15/SPIN2003	188	984	0.0	6436	1.7	699	0.2	–	0.1
16/STACK (NON-BL.)	352	2550	MO	34046	133.7	18603	128.6	8249	12.5
17/STACK	648	3626	MO	35500	38.7	7616	20.2	2723	2.3
18/BOOP	7488	25929	0.0	1446	7.8	10776	361.1	–	0.1
19/MOZILLA-VUL.-FIXED	1648	8050	0.0	77053	84.2	3723	4.3	–	1.7
20/PETERSON	2048	8988	0.0	22951	15.5	2373	2.3	–	1.2
21/SZYMANSKI	8448	35896	0.1	–	TO	9597	35.8	–	11.0

of our new approach, the oracle can be deactivated, turning BREACH into the refined version of the classical backward search (Algorithm 1). Due to efficiency limitations of the underlying data structures, we do not add candidate vertices that involve two threads or more (which we found to be a good trade-off between efficiency and proof minimization). To apply BREACH to the C programs, we extended the abstract language interface of the C software model checker SATABS to TTS. SATABS implements the CEGAR loop based on a symmetry-aware predicate-abstraction technique [10], and handles function calls by inlining. All experiments are performed on a 3GHz Intel Xeon machine with 20 GB memory, running 64-bit Linux, with a timeout of 30 minutes.

Evaluation. Table 1 presents results for various configurations of our implementation. Columns on the left show the benchmark id and name, and the total number of thread states and transitions emerged in the last, and always most expensive, CEGAR iteration. Remaining columns show details for:

GKM: Our coverability oracle (stand-alone);

BC: Refined version of the classical backward algorithm (Algorithm 1);

MCOV: Our MCOV algorithm (Algorithm 3);

MCOV/GKM: The MCOV algorithm equipped with the coverability oracle GKM.

For each approach we show the total model checking run time, and in addition for backward-directed algorithms the number of iterations.

The results demonstrate that our new approach outperforms the classical algorithms: MCOV/GKM solves *all* 21 programs, and MCOV 17 instances, compared to 13 and 9 for the classical backward algorithm and the coverability oracle, respectively. Comparing the results for BC and MCOV clearly shows that the uncoverability proofs the latter




Fig. 3. Cactus plot on the 21 multi-threaded C programs comparing our MCOV and MCOV/GKM approaches to various existing ones; due to broadcasts, the limit for tools marked with * is $k = 17$

generates are much smaller. This is reflected by a strict decrease of the iteration count in 17 cases. In the majority of cases, this improvement manifests in the running time: MCOV outperforms BC on 13 programs (often significantly), compared to 4 the other way around. Furthermore, the results for MCOV/GKM show that the coverability oracle can substantially reduce the cost of unhelpful candidates, showing their synergies (observe that running GKM and BC stand-alone in parallel is not helpful). As a result, the positive effect is amplified: compared to BC, the iteration count strictly decreases on *all* programs.

To measure the difference between standard and minimal uncoverability proofs, we removed the bound on candidate vertices (in return for longer runtimes). In this setup, we observed the following reductions (averaged): the the longest traversed path drops from 28 to 14 (-50%), the threads included in the proof from 6 to 2 (-67%), and the proof size in terms of minimal states from 22518 to 1222 (-95%). While the classical backward-approach includes up to eight threads in a proof, our approach always generates minimal uncoverability proofs which involve no more than *two threads*. With the bound on candidate vertices mentioned above and used for Table 1, the reductions are only marginally smaller (e.g., the previous thread number increases by one).

Comparison. There exist a number of other approaches to the coverability problem. We compare to the following tools (all available online):

IST-BC: Classical backward search using interval sharing-trees (v1.0.3) [18];

PETR-BC: Refined backward search with structural invariants (v0.1) [29,30];

TINA-KM: Karp-Miller procedure (v3.0.0) [5];

CSC-KM: Refined Karp-Miller procedure using interval sharing-trees (v0.1) [22];

EEC-AR: Pure forward algorithm with enumerative refinement (v1.0.3) [21];

TSI-AR: Variant of [21] using backward underapprox. for refinement (v1.0.3) [19].

Only IST-BC and PETR-BC support broadcast primitives. In order to allow for a meaningful comparison, we translated abstract TTS templates generated by SATABS into (transfer) Petri nets and replaced the model checker back-end.

Figure 3 depicts total model checking run times (scaled logarithmically) for all methods as “cactus plot”: the horizontal axis represents the number of programs the respective method could successfully handle, and the vertical axis the time needed to solve this number if they were ran in parallel. The results demonstrate significant improvements over all previous methods: only MCOV/GKM is able to solve all 21 programs, followed by MCOV stand-alone (17), PETR-BC (15), BC (13), EEC-AR (11), TINA-KM and CSC-KM (9), IST-BC (7), and TSI-AR (1).

The improvement over the best previous approach (PETR-BC) shows that our new approach is able to guide the search more effectively than structural invariant heuristics, which are known to often yield invariants that are irrelevant to the safety property or too imprecise [15]. The inferior performance of our underlying classical backward algorithm (BC) to PETR-BC indicates that the observed improvements are not just owed to clever implementation, but rather the result of our novel approach.

6 Related Work

Algorithmic solutions to coverability analysis were first proposed for *vector addition systems* in a landmark paper by Karp and Miller [26]. The solution constructs a pseudo-reachability tree by forward exploration and replaces newly discovered states that are strictly greater than predecessors by their limit. It has a non-primitive recursive worst-case complexity [32]. The purpose there was mainly to show decidability of the coverability problem for VASes and the equivalent Petri nets. The technique is implemented in the tool TINA-KM [5]. It cannot be extended to broadcast primitives [12]. An improvement of this procedure that computes minimal coverability sets is [22].

To afford more flexibility in modeling parametrized programs, various algorithms were later proposed for WSTS, originally in a pure backward fashion [2], which was implemented in the tools IST-BC and PETR-BC [29], later as forward exploration [14,34]. The paradigm presented in [21] (and implemented in the tool EEC-AR) is also a pure forward algorithm; it constructs abstractions of increasing precision. In contrast to the paradigm of EEC-AR, the implementation itself does not support broadcasts. Other approaches are the backward and forward unfolding algorithms from [3] and [25].

Solutions combining forward and backward exploration are rare; we are only aware of the methods described in [15] and [19]. The authors of [15] propose to use a CSC-KM-like approach to compute overapproximations of the coverability set, which are then used in a *subsequent* backward exploration to prune the search space. Our experimental results demonstrate, however, that this approach cannot cope with programs of the sizes we consider. In [19], the authors combine overapproximations computed in a forward fashion, which are refined by using backward underapproximations; the approach is implemented in the tool TSI-AR. On an abstract level, our algorithm can be seen as the dual of this approach. To the best of our knowledge, our approach is the first to combine forward propagation of *underapproximations* with backward propagation of *overapproximations* to the coverability problem in WSTS.

7 Conclusion

We introduced a new approach to the coverability problem in WSTS. The novelty of our algorithm is the way it proves uncoverable instances via a sequence of many inexpensive uncoverability proofs. Our algorithm can be used to check assertion failures, mutual exclusion violations and many other properties for parametrized programs communicating via mutexes, shared variables or common concurrency primitives such as broadcasts.

We demonstrated in extensive experiments on large benchmarks, generated by the software model checker SATABS from C programs, that our algorithm outperforms the best known coverability approach by orders of magnitude, enabling the verification of programs which are out of scope of the previous technology. The experiments also reveal that our approach is able to guide the search far more effectively than existing structural invariant heuristics [13,8]. We conclude from our experiments that programs tend to feature minimal uncoverability proofs with fewer and smaller minimal elements compared to those targeted by existing methods.

The ideas we have presented, supported by the simplicity of their implementation, are naturally applicable to coverability methods in general. We believe, for example, that while our method outperforms techniques based on structural net invariants, even more practical benefit is achievable by combining these strategies.

Acknowledgments. We wish to thank Michael Tautschnig for assistance with SATABS, and Pierre Ganty, Leopold Haller, Philipp Rümmer and Emelie Vollmer for their insightful comments on earlier drafts of this work.

References

1. Abdulla, P.A.: Well (and better) quasi-ordered transition systems. *Bulletin of Symbolic Logic* 16(4) (2010)
2. Abdulla, P.A., Cerans, K., Jonsson, B., Yih-Kuen, T.: General decidability theorems of infinite-state systems. In: *Logic in Computer Science, LICS* (1996)
3. Abdulla, P.A., Iyer, S.P., Nylén, A.: SAT-solving the coverability problem for Petri nets. *Formal Methods in System Design, FMSD* (2004)
4. Ball, T., Rajamani, S.: The SLAM project: debugging system software via static analysis. In: *Principles of Programming Languages, POPL* (2002)
5. Berthomieu, B., Vernadat, F.: The Tina tool, release 2.9.6, LAAS/CNRS (November 2009), <http://homepages.laas.fr/bernard/tina/>
6. Cardoza, E., Lipton, R.J., Meyer, A.R.: Exponential space complete problems for Petri nets and commutative semigroups: Preliminary report. In: *STOC*, pp. 50–54 (1976)
7. Clarke, E., Kroning, D., Sharygina, N., Yorav, K.: SATABS: SAT-Based Predicate Abstraction for ANSI-C. In: Halbwachs, N., Zuck, L.D. (eds.) *TACAS 2005*. LNCS, vol. 3440, pp. 570–574. Springer, Heidelberg (2005)
8. Delzanno, G., Raskin, J.-F., Van Begin, L.: Attacking Symbolic State Explosion. In: Berry, G., Comon, H., Finkel, A. (eds.) *CAV 2001*. LNCS, vol. 2102, pp. 298–310. Springer, Heidelberg (2001)

9. Delzanno, G., Raskin, J.-F., Van Begin, L.: Towards the Automated Verification of Multi-threaded Java Programs. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, pp. 173–187. Springer, Heidelberg (2002)
10. Donaldson, A., Kaiser, A., Kroening, D., Wahl, T.: Symmetry-Aware Predicate Abstraction for Shared-Variable Concurrent Programs. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 356–371. Springer, Heidelberg (2011)
11. Emerson, A., Namjoshi, K.K.: On model checking for non-deterministic infinite-state systems. In: Logic in Computer Science (LICS), pp. 70–80 (1998)
12. Esparza, J., Finkel, A., Mayr, R.: On the verification of broadcast protocols. In: Logic in Computer Science, LICS (1999)
13. Esparza, J., Melzer, S.: Verification of safety properties using integer programming: Beyond the state equation. Formal Methods in System Design, FMSD (2000)
14. Finkel, A., Goubault-Larrecq, J.: Forward Analysis for WSTS, Part II: Complete WSTS. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S., Thomas, W. (eds.) ICALP 2009. LNCS, vol. 5556, pp. 188–199. Springer, Heidelberg (2009)
15. Finkel, A., Raskin, J.-F., Samuelides, M., Begin, L.V.: Monotonic extensions of Petri nets: Forward and backward search revisited. ENTCS (2002)
16. Finkel, A., Schnoebelen, P.: Well-structured transition systems everywhere! Theoretical Computer Science, TCS (2001)
17. Flanagan, C., Qadeer, S.: Thread-Modular Model Checking. In: Ball, T., Rajamani, S.K. (eds.) SPIN 2003. LNCS, vol. 2648, pp. 213–224. Springer, Heidelberg (2003)
18. Ganty, P., Meuter, C., Delzanno, G., Kalyon, G., Raskin, J.-F., Van Begin, L.: Symbolic data structure for sets of k-uples. Technical report, Université Libre de Bruxelles (2007)
19. Ganty, P., Raskin, J.-F., Van Begin, L.: A Complete Abstract Interpretation Framework for Coverability Properties of WSTS. In: Emerson, E.A., Namjoshi, K.S. (eds.) VMCAI 2006. LNCS, vol. 3855, pp. 49–64. Springer, Heidelberg (2005)
20. Ganty, P., Raskin, J.-F., Van Begin, L.: From Many Places to Few: Automatic Abstraction Refinement for Petri Nets. In: Kleijn, J., Yakovlev, A. (eds.) ICATPN 2007. LNCS, vol. 4546, pp. 124–143. Springer, Heidelberg (2007)
21. Geeraerts, G., Raskin, J.-F., Begin, L.V.: Expand, enlarge and check: New algorithms for the coverability problem of WSTS. JCSS (2006)
22. Geeraerts, G., Raskin, J.-F., Van Begin, L.: On the Efficient Computation of the Minimal Coverability Set for Petri Nets. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) ATVA 2007. LNCS, vol. 4762, pp. 98–113. Springer, Heidelberg (2007)
23. German, S., Sistla, P.: Reasoning about systems with many processes. Journal of the ACM, JACM (1992)
24. Gupta, A., Popeea, C., Rybalchenko, A.: Threader: A Constraint-Based Verifier for Multi-threaded Programs. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 412–417. Springer, Heidelberg (2011)
25. Kaiser, A., Kroening, D., Wahl, T.: Dynamic Cutoff Detection in Parameterized Concurrent Programs. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 645–659. Springer, Heidelberg (2010)
26. Karp, R.M., Miller, R.E.: Parallel program schemata. J. Comput. Syst. Sci. 3(2), 147–195 (1969)
27. Lu, S., Park, S., Seo, E., Zhou, Y.: Learning from mistakes: a comprehensive study on real world concurrency bug characteristics. In: Architectural Support for Programming Languages and Operating Systems, ASPLOS (2008)
28. Malkis, A., Podelski, A., Rybalchenko, A.: Thread-Modular Verification Is Cartesian Abstract Interpretation. In: Barkaoui, K., Cavalcanti, A., Cerone, A. (eds.) ICTAC 2006. LNCS, vol. 4281, pp. 183–197. Springer, Heidelberg (2006)

29. Meyer, R., Strazny, T.: Petruccio: From Dynamic Networks to Nets. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 175–179. Springer, Heidelberg (2010)
30. Meyer, R., Strazny, T.: An algorithmic framework for coverability in well-structured systems. In: Application of Concurrency to System Design, ACS D (2012)
31. Peierls, T., Goetz, B., Bloch, J., Bowbeer, J., Lea, D., Holmes, D.: Java Concurrency in Practice. Addison-Wesley Professional (2005)
32. Rackoff, C.: The covering and boundedness problems for vector addition systems. Theoretical Computer Science, TCS (1978)
33. Schnoebelen, P.: Revisiting Ackermann-Hardness for Lossy Counter Machines and Reset Petri Nets. In: Hliněný, P., Kučera, A. (eds.) MFCS 2010. LNCS, vol. 6281, pp. 616–628. Springer, Heidelberg (2010)
34. Zufferey, D., Wies, T., Henzinger, T.A.: Ideal Abstractions for Well-Structured Transition Systems. In: Kuncak, V., Rybalchenko, A. (eds.) VMCAI 2012. LNCS, vol. 7148, pp. 445–460. Springer, Heidelberg (2012)