REGULAR PAPER

An abstraction refinement approach combining precise and approximated techniques

**Natasha Sharygina · Stefano Tonetta ·
Aliaksei Tsitovich**

Abstract Predicate abstraction is a powerful technique to reduce the state space of a program to a finite and affordable number of states. It produces a conservative over-approximation where concrete states are grouped together according to a given set of predicates. A precise abstraction contains the minimal set of transitions with regard to the predicates, but as a result is computationally expensive. Most model checkers therefore approximate the abstraction to alleviate the computation of the abstract system by trading off precision with cost. However, approximation results in a higher number of refinement iterations, since it can produce more false counterexamples than its precise counterpart. The refinement loop can become prohibitively expensive for large programs. This paper proposes a new approach that employs both precise (slow) and approximated (fast) abstraction techniques within one abstraction-refinement loop. It allows computing the abstraction quickly, but keeps it precise enough to avoid too many refinement iterations. We implemented the new algorithm in a state-of-the-art software model checker. Our tests with various real-life benchmarks show that the new approach almost systematically outperforms both precise and imprecise techniques.

Keywords Predicate abstraction · Precise abstraction ·
Approximated abstraction · CEGAR

N. Sharygina · A. Tsitovich (✉)
University of Lugano, Lugano, Switzerland
e-mail: aliaksei.tsitovich@usi.ch

S. Tonetta
Fondazione Bruno Kessler, Trento, Italy

N. Sharygina
School of Computer Science, Carnegie Mellon University,
Pittsburgh, USA

1 Introduction

Predicate abstraction [16,20], when combined with reachability analysis and an automated abstraction refinement mechanism (also known as *Counterexample Guided Abstraction Refinement* (CEGAR)[5,12]), is an effective model checking strategy. The CEGAR-based verification consists of constructing and evaluating a finite-state system that is an abstract model of the original system with respect to a set of predicates.

The abstract model is a conservative over-approximation of the original program with respect to the set of given predicates. Thus, if the property holds on the abstract model, it also holds on the original program. The drawback of the conservative abstraction is that when model checking of the abstract program fails, it may produce a counterexample that does not correspond to any concrete counterexample. This is called a *spurious counterexample*. When a spurious counterexample is encountered, *refinement* is performed by adjusting the set of predicates in a way that eliminates the given counterexample. The overall efficiency of verification is highly dependent on the efficiency of the abstraction and refinement procedures.

Computing the abstract model relies on enumerating the abstract states and checking, for each pair of states, if there exists an abstract transition. This computation is expensive since it requires an exponential number of calls to a theorem prover [2,3,18]. In [10,29,30], the abstraction is computed by means of dedicated decision procedures based on BDDs, SAT or SAT modulo theories (SMT). As another direction, various techniques have been proposed to alleviate this computation by approximating the abstract transition relation (see, for example, [2–4,17,26]).

We distinguish between *precise* abstraction and *approximated* abstraction (as also done, for example, in [12,17,26]):

a precise abstraction is minimal in the sense that it contains only those transitions that correspond to some transition in the concrete model; instead, an approximated abstraction is a further over-approximation of the minimal abstract model so that the transition relation is relaxed. In the paper, we will refer to the latter simply as approximation.

Approximation techniques are important because they allow a less expensive (as compared to precise abstraction) computation of the abstract transition relation. Cartesian abstraction [4], for example, loses every relationship among predicates, but has been successfully used to verify large programs, such as operating system device drivers. However, abstraction approximations add spurious behaviors in addition to the spurious counterexamples resulting from precise abstraction. In order to rule out this kind of "impurity", the approximation must be refined without changing the set of predicates and focused only on the spurious transitions caused by the approximation [17]. This procedure on its own might become very costly and does not scale to verification of large programs.

When refining the abstract model, we distinguish between two types of spurious behavior (as also done in [14]). (1) *Spurious path* is due to the over-approximating nature of the precise abstraction: states are merged together so that some resulting paths cannot be simulated on the concrete system. This happens when the set of predicates is not sufficient to capture the relevant behaviors of the concrete system. (2) *Spurious transitions* are abstract transitions which do not have corresponding concrete transitions. By definition, spurious transitions cannot appear in the most precise abstraction and are caused by using the approximation techniques. Clearly, the efficiency of the approximated abstraction depends on a tradeoff between time spent in computing the abstraction and refining spurious transitions.

In order to illustrate the abstraction approximation and its refinement procedures, consider the example of Fig. 1a. The variable x is assigned non-deterministically with an unknown value "∗". The property we verify is the reachability of line l3. It never can be reached since the condition !(x<y) at line l2 never holds (together with the guard x<0 at line l1, which is necessary to avoid integer overflow). Thus, if in the abstract program there is a path leading to the assertion, then it is spurious. The predicates x<0 and x<y are sufficient to prove the property. However, approximate methods like Cartesian abstraction cannot prove it because they cannot infer that after the assignment y=x+1, the condition (!(x<0) || !(x<y)) is true. Thus, most model checkers that use such abstractions refine the transition relation by adding a constraint that removes the spurious transition.

In order to experience the difference in performance between precise and approximated abstractions, let us extend the previous example in order to have more spurious behaviors. The program of Fig. 1b has one more variable and a

Fig. 1 Sample program for which the approximated abstraction causes spurious transitions

```
void main() {
  int x=*;
  int y;
l0: y=x+1;
l1: if (x<0)
l2: if (!(x<y))
l3: assert(0);
l4: }
        (a)

void main() {
  int x = *;
  int y;
  int z;
  y=x+1;
  z=y-1;
  if (x<0||y>0||z>0){
  x++;
  z++;
  if (y<z && z>x)
  assert(0); } }
        (b)
```

slightly more complex control flow graph. As before the assertion is not reachable, and all abstract counterexamples are spurious. Though, if we consider the predicates in the guards of the program, an approximated abstraction may produce many spurious behaviors. Table 1 reports the verification results obtained with the SATABS model checker [15], by running approximated and precise abstractions. The final number of predicates is in all cases 10. The approximated abstraction spends most of time in refining the transition relation (Ref). Since it runs for 12 iterations (or even 42 in case when we used the refinement procedure of [17]), also the time for the verification (MC) is not negligible. On the contrary, the precise abstraction takes only 2 iterations to terminate (the first refinement is necessary to add a sufficient set of predicates). Nevertheless, the amount of time spent in computing the abstraction is too high for such example.

A low number of refinement iterations is fundamental for the success of the CEGAR loop, especially when applied to industrial benchmarks: in fact, when the system is complex, the number of predicates required to verify the property becomes high, and the time spent in the reachability (model checking) procedure grows exponentially. For this reason, it is of paramount importance to avoid as many redundant iterations as possible: even a single saved iteration can result in a huge saving in time for large systems.

Contributions This paper presents a CEGAR-based technique that controls the number of iterations and reduces the verification time by interleaving precise (but slow) and approximated (but fast) abstractions. The abstraction is first computed with a high level of approximation exploiting the weakest precondition of the predicates. Then, during the refinement step, our technique uses the SAT-based quantifier elimination in order to compute a precise abstraction. We also show how precise component computation can be

Table 1 Verification results on the example of Fig. 1b

	Total	Abs	MC	Ref	Iter
Approximated abstraction [17]	5.817	0.063	2.659	2.112	42
Approximated abstraction [24]	1.469	0.046	0.501	0.617	12
Precise abstraction	3.591	3.478	0.076	0.01	2
New approach	0.467	0.039	0.161	0.189	4

Total, Abs, MC, Ref refer to the time, in seconds, for total verification, abstraction, model checking and refinement respectively; *Iter* refers to the number of iterations of the abstraction-refinement loop

heuristically limited in order to avoid possible exponential blow ups.

The difficulty that we would experience in computing the precise abstraction of the whole program is avoided by exploiting the localized abstraction: as in static analysis [33], in most model checkers (such as SLAM [3], BLAST [23], SATABS [15], F-Soft [25]) the abstract model keeps the control flow graph of the original program and has a different abstract transition relation for each location of the control-flow graph. [1] This way, during the refinement step, we add the constraints built with a precise abstraction only to relevant transition relations, affecting only those parts of the system that caused the spurious counterexample.

In order to illustrate the immediate advantages of our approach, consider the fourth line of Table 1 that is based on the implementation of our technique. Our approach is able to avoid both a high number of iterations and an expensive abstraction, resulting in an optimized verification time.

We performed a thorough evaluation comparing the new technique with the purely precise and imprecise counterparts. Our tests with various real-life benchmarks show a systematic advantage of our approach over both precise and imprecise techniques reaching up to 90% improvement in time.

Overall, the new technique manages the verification complexity by using the precise abstraction on demand and locally. The advantage is that the expensive abstraction is only used on a small portion of the program, yet the higher quality of abstraction refinement is sufficient to reduce the number of refinement iterations, thus improving the overall performance.

Related work

This paper addresses the problem of refining the abstraction in the presence of spurious transitions. The solution was first given by Das and Dill [17] whose technique consists of removing one spurious transition at every refinement iteration. The approach may be very expensive because it requires a high number of iterations of the abstraction–refinement

loop. In practice, the technique is not feasible for real systems.

Many works such as [2] improved the refinement by strengthening the condition added to the transition relation to remove more spurious transitions. The idea in [2] is to syntactically simplify the condition and to check if a larger set of spurious transitions is found.

In [11,24,25], a different technique is presented based on SAT techniques. Transitions are simulated over the concrete program by means of SAT formulas. If the transition is not concretizable the SAT solver will produce a resolution proof of the unsatisfiability. It is then possible to extract from the proof either a core set of predicates or a constraint sufficient to remove the spurious transition. Though, in principle, the technique can remove many spurious transitions at once, the efficiency strongly depends on the unsatisfiability proof. In the worst case, it may require a number of abstraction refinements exponential in the number of predicates.

The technique of [26] also exploits the unsatisfiability proof but it is based on interpolation. The interpolant produced by the proof is indeed an over-approximation of the exact abstraction able to remove the spurious transition. As in the case of unsat cores, the technique depends on the heuristics to produce unsatisfiability proofs. The interpolant is not always enough strong to remove all spurious transitions.

This paper instead proposes a greedy approach where all spurious transitions between two locations are removed. The idea is that the computation can be efficient because it is localized and on-demand. The technique inherits the efficiency of the approximated abstraction which is used any time new predicates are discovered. At the same time, the precision of the minimal abstraction is exploited whenever spurious transitions are found.

Summary

The paper is organized as follows: Sect. 2 gives an overview of related abstraction refinement techniques; Sect. 3 describes our new approach; Sect. 4 presents the experimental evaluation; finally, Sect. 5 draws the conclusions.

---

[1] Localized abstraction is further investigated in [22,23].

2 Background

2.1 Transition systems

We consider programs as Transition Systems. TSs are defined by a set V of state variables. We use V' to denote the set of next state variables $\{v'\}_{v \in V}$, where v' represents the next value of v. The set S_V of states is given by all assignments to the variables V. Given a state s, s' denotes the corresponding assignment to the next state variables, i.e. $s' = s[V'/V]$. Transitions are represented as pairs of states. For each transition $t = (s_1, s_2)$, we use $in(t)$ and $out(t)$ to denote resp. s_1 and s_2. Given a formula ϕ, we write $\phi[V/V']$ to denote the result of substituting every free occurrence of every variable $v' \in V'$ with its corresponding v. We use $\exists V(\phi)$ to denote the existential quantification of every variable in V.

Definition 1 A Transition System (TS) is a tuple $M = \langle V, I, T \rangle$, where

- V is a set of variables;
- $I(V)$ is a formula that represents the initial states;
- $T(V, V')$ is a formula that represents the transitions.

A state s is initial iff $s \models I(V)$. Given two states s_1 and s_2, there exists a transition t between s_1 and s_2 iff $s_1, s_2' \models T(V, V')$. A path of M is a finite sequence π of transitions $t_0, t_1, \ldots, t_n$ such that $in(t_0) \models I$, and, for every $0 \le i < n$, $out(t_i) = in(t_{i+1})$. In general, given a transition relation T, we use $\pi \models T$ to denote that $\pi[i] \models T$ for every $0 \le i \le |\pi|$.

Example 1 Consider the program of Fig. 1(a). It can be represented by the TS $M = \langle V, I, T \rangle$, where

- $V := \{x, y, pc\}$, where pc is the program counter;
- $I := (pc = l_0)$;
- $T := (pc = l_0) \rightarrow (pc' = l_1 \wedge y' = x + 1 \wedge x' = x) \wedge$
  $(pc = l_1 \wedge x < 0) \rightarrow (pc' = l_2 \wedge x' = x \wedge y' = y) \wedge$
  $(pc = l_1 \wedge !x < 0) \rightarrow (pc' = l_4 \wedge x' = x \wedge y' = y) \wedge$
  $(pc = l_2 \wedge !x < y) \rightarrow (pc' = l_3 \wedge x' = x \wedge y' = y) \wedge$
  $(pc = l_2 \wedge x < y) \rightarrow (pc' = l_4 \wedge x' = x \wedge y' = y) \wedge$
  $(pc = l_3) \rightarrow (pc' = l_4 \wedge x' = x \wedge y' = y)$

2.2 Abstraction

Definition 2 Given two TSs $M = \langle V, I, T \rangle$ and $\hat{M} = \langle \hat{V}, \hat{I}, \hat{T} \rangle$, a relation $H(V, \hat{V})$ is an *abstraction relation* [13] iff the following conditions hold:

- every initial state of M corresponds to an initial state of $\hat{M}$; namely, if $s \models I(V)$, then there exists a state $\hat{s}$ of $\hat{M}$ such that $\hat{s} \models \hat{I}(\hat{V})$ and $s, \hat{s} \models H(V, \hat{V})$;

- every transition of M corresponds to a transition of $\hat{M}$; namely, if $s_1, \hat{s}_1 \models H(V, \hat{V})$, and $s_1, s_2' \models T(V, V')$, then there exists a state $\hat{s}_2$ of $\hat{M}$ such that $s_2, \hat{s}_2 \models H(V, \hat{V})$ and $\hat{s}_1, \hat{s}_2' \models \hat{T}(V, V')$.

If such relation exists, we say that $\hat{M}$ is an *abstraction* of M, or M *refines* $\hat{M}$ ($M \preceq \hat{M}$).

Definition 3 Given the abstraction relation H, we define the *abstraction function* $\alpha_H : 2^{S_V} \rightarrow 2^{S_{\hat{V}}}$ and the *concretization function* $\gamma_H : 2^{S_{\hat{V}}} \rightarrow 2^{S_V}$ as follows:

- $\alpha_H(Q) = \{\hat{s} \in S_{\hat{V}} \mid \text{ there exists } s \in Q \text{ s.t. } s, \hat{s} \models H(V, \hat{V})\}$, for every $Q \subseteq S_V$;
- $\gamma_H(\hat{Q}) = \{s \in S_V \mid \text{ there exists } \hat{s} \in \hat{Q} \text{ s.t. } s, \hat{s} \models H(V, \hat{V})\}$, for every $\hat{Q} \subseteq S_{\hat{V}}$.

We extend γ to transitions and paths so that

- $\gamma_H(\hat{t}) = \{t \mid in(t) \in \gamma(in(\hat{t})), out(t) \in \gamma(out(\hat{t}))\}$, for every transition $\hat{t}$ of $\hat{M}$.
- $\gamma_H(\hat{\pi}) = \{\pi \mid \pi[i] \in \gamma(\hat{\pi}[i]) \text{ for every } 0 \le i \le |\hat{\pi}|\}$, for every path $\hat{\pi}$ of $\hat{M}$.

If F is a subset of Q, F is an invariant for a system M iff for all paths of M all states of the paths belong to F. The abstraction relation we defined preserves invariants (and more in general all universal properties in $\forall CTL^*$ [13]), so that if $M \preceq \hat{M}$, and $\alpha_H(F)$ is an invariant of $\hat{M}$, then F is an invariant of M (though, in general, the reverse does not hold). Given a TS $M = \langle V, I, T \rangle$, an abstraction $\hat{M} = \langle \hat{V}, \hat{I}, \hat{T} \rangle$ of M is said to be *precise* when every abstract initial state and transition of $\hat{M}$ corresponds, respectively, to a concrete initial state and transition of M. Given the abstraction relation H, $\hat{M}$ can be obtained as

- $\hat{I}_H(\hat{V}) = \exists V(I(V) \wedge H(V, \hat{V}))$,
- $\hat{T}_H(\hat{V}, \hat{V}') = \exists V \exists V'(T(V, V') \wedge H(V, \hat{V}) \wedge H(V', \hat{V}'))$

The precise abstraction is also called *minimal* or *existential* or *exact* or *eager* abstraction [13].

Given a TS $M = \langle V, I, T \rangle$, let P be a set of predicates and $\hat{v}_p$ an abstract variable for every predicate $p \in P$. The set of abstract variables is the set $\hat{V}_P = \{\hat{v}_p\}_{p \in P}$. The abstraction relation for predicate abstraction is defined as follows:

$$H_P(V, \hat{V}_P) = \bigwedge_{p \in P} \hat{v}_p \leftrightarrow p(V) \tag{1}$$

The minimal predicate abstraction is the TS $\hat{M} = \langle \hat{V}_P, \hat{I}_P, \hat{T}_P \rangle$, where:

- $\hat{I}_P(\hat{V}_P) = \exists V(I(V) \wedge \bigwedge_{p \in P} \hat{v}_p \leftrightarrow p(V))$

$-\quad \hat{T}_P(\hat{V}_P, \hat{V}'_P) = \exists V \exists V'(T(V, V') \wedge \bigwedge_{p \in P}(\hat{v}_p \leftrightarrow p(V) \wedge \hat{v}'_p \leftrightarrow p(V')))$.

2.2.1 Quantifier elimination

In order to model check the abstract TS, it is necessary to compute the set of successors of abstract states. This requires the removal of the quantifiers from the definition of the abstract transition relation. In general, given a transition relation T and a set of predicates P, *to compute $\hat{T}_P$* means to find a quantifier-free formula that is equivalent to $\hat{T}_P$.

Example 2 Consider the TS described in the Example 1 and the predicates $P_1 := (x < 0)$ and $P_2 := (x < y)$. Let the abstract variables $\hat{v}_1$ and $\hat{v}_2$ correspond, respectively, to P_1 and P_2. We do not abstract the program counter. The abstract transition relation results to be equivalent to

$-\quad \hat{T}_P \equiv (pc = l_0 \wedge \hat{v}_1) \rightarrow (pc' = l_1 \wedge !\hat{v}'_2) \wedge$
$(pc = l_0 \wedge !\hat{v}_1) \rightarrow (pc' = l_1) \wedge$
$(pc = l_1 \wedge \hat{v}_1) \rightarrow (pc' = l_2 \wedge \hat{v}'_1 = \hat{v}_1 \wedge \hat{v}'_2 = \hat{v}_2) \wedge$
$(pc = l_1 \wedge !\hat{v}_1) \rightarrow (pc' = l_4 \wedge \hat{v}'_1 = \hat{v}_1 \wedge \hat{v}'_2 = \hat{v}_2) \wedge$
$(pc = l_2 \wedge !\hat{v}_2) \rightarrow (pc' = l_3 \wedge \hat{v}'_1 = \hat{v}_1 \wedge \hat{v}'_2 = \hat{v}_2) \wedge$
$(pc = l_2 \wedge \hat{v}_2) \rightarrow (pc' = l_4 \wedge \hat{v}'_1 = \hat{v}_1 \wedge \hat{v}'_2 = \hat{v}_2) \wedge$
$(pc = l_3) \rightarrow (pc' = l_4 \wedge \hat{v}'_1 = \hat{v}_1 \wedge \hat{v}'_2 = \hat{v}_2)$

In hardware and software verification, different techniques have been conceived to compute $\hat{T}_P$. In symbolic model checking [9] of finite state machines, the existential quantification can be removed either by a Shannon expansion technique when using BDDs [8] or by SAT techniques when using CNF [32]. In software model checking, the problem is exacerbated by the fact that the concrete transition relation may contain first-order terms. The abstract transition relation can be obtained by enumerating the abstract states, and checking if, for each pair of states, there exists an abstract transition. As it is done by most software model checkers, this requires an exponential number of calls to a theorem prover [3,18]. In [15], a SAT solver is exploited to find all possible solutions. We refer to this technique as *SATQE*.

2.3 Abstraction approximation

Precise abstractions are very expensive to compute because of the existential quantification operations. Thus, in practice, model checkers use approximations to trade-off precision with complexity.

Definition 4 Formally, given $M_H = \langle V, I_H, T_H \rangle$ and $\tilde{M} = \langle V, \tilde{I}, \tilde{T} \rangle$, we say that $\tilde{M}$ is an *approximation* of M_H ($M_H \precsim \tilde{M}$) iff the following formulas are valid:

$-\quad I_H \rightarrow \tilde{I}$, i.e., every initial state of the minimal abstraction is an initial state in the approximation;

$-\quad T_H \rightarrow \tilde{T}$, i.e., every transition of the minimal abstraction is a transition in the approximation.

Intuitively, $\tilde{M}$ has more initial states and transitions than M_H. Note that an approximation is also an abstraction namely, if $M_H \precsim \tilde{M}$, then $M_H \preceq \tilde{M}$. However, the set of predicates is not affected, in the sense that $\tilde{M}$ and M_H have the same abstract variables.

2.3.1 Approximation techniques

Many approximation techniques have been developed both in hardware and software verification. Their aim is to alleviate the computation of $\hat{T}_P$. The easiest way is to reduce the scope of quantifiers. This can be done with *early quantification* [13], by pushing quantifiers in front of predicates. *Predicate partitioning* [24] approximates $\hat{T}_P$ by taking the conjunction of its projections over subsets of predicates. This technique is pushed to its limit by Cartesian abstraction [4] that, given a set of states Q, approximates transition relation with the product of the projections on each variable. This way, the approximated abstraction ignores every relation among predicates.

2.4 Spurious behaviors

The over-approximation nature of the abstraction as we define may generate spurious paths even in the case of precise abstraction. Spurious paths are sequences of transitions that satisfy the abstract transition relation, but not the concrete one.

Definition 5 (*Spurious path*) Given a TS $M = \langle V, I, T \rangle$, an abstraction $\hat{M} = \langle \hat{V}, \hat{I}, \hat{T} \rangle$, and a sequence $\hat{\pi}$ of transitions of $\hat{M}$, we say that $\hat{\pi}$ is a spurious path iff $\hat{\pi} \models \hat{T}$ and $\pi \not\models T$ for every $\pi \in \gamma(\hat{\pi})$.

In order to refine the abstraction and remove a spurious path, refinement procedures need to add more predicates to the abstraction. There are different techniques to discover the new set of predicates, either based on weakest precondition [6], interpolation [22], or UNSAT core [21].

Besides spurious path, approximated abstraction generates another kind of spurious behavior, called spurious transitions. Spurious transitions are transitions that satisfy the abstract transition relation, but not the concrete one.

Definition 6 (*Spurious transition*) Given a TS $M = \langle V, I, T \rangle$, an abstraction $\hat{M} = \langle \hat{V}, \hat{I}, \hat{T} \rangle$, and a transition $\hat{t}$ of $\hat{M}$, we say that $\hat{t}$ is a spurious transition iff $\hat{t} \models \hat{T}$ and $t \not\models T$ for every $t \in \gamma(\hat{t})$.

In order to refine an approximation that contains a spurious transition, a new transition relation is obtained by adding

a constraint in conjunction to the old abstract transition relation. As a result, the spurious counterexample is ruled out. Different techniques use as such constraint either the exact encoding of the spurious transition [17], or the UNSAT core produced by the SAT solver when checking if the transition is spurious [24], or an interpolant between the exact abstraction and the current approximated abstraction [26].

3 The synergy algorithm

This section proposes a new refinement algorithm. It uses both the fast and precise types of abstraction to gain verification efficiency. It is independent of any particular technique used to define either procedure.

The algorithm implements the standard CEGAR loop. Each iteration of the CEGAR loop is composed of an abstraction step, a model checking step, a simulation step, and finally a refinement step.

We first present the high-level overview of the combined algorithm and then describe the specifics of the new refinement procedures. For simplicity, we first present the algorithm with regard to a monolithic transition relation. In Sect. 3.3 we extend it to the case where a transition relation is defined for every location of the program.

The algorithm is parameterized by a number of subroutines that take care of the abstraction and refinement. In particular, the algorithm contains the following procedures:

- `FastAbstraction`: given a set of predicates Π and a concrete transition relation T, it computes an over-approximation of $\hat{T}_\Pi$.
- `PreciseAbstraction`: given a set of predicates Π and a concrete transition relation T, it computes the minimal abstraction $\hat{T}_\Pi$.
- `SpuriousTransition`: given a path π in $\hat{M}$, it returns a function σ_{ST} that maps every transition t in π to a set of predicates P, s.t., $P \subseteq \Pi$ and $t \not\models \hat{T}_P$.
- `SpuriousPath`: given a path π in $\hat{M}$, it returns a function σ_{SP} that maps every transition t in π to a set of predicates P, s.t. $\pi \not\models \hat{T}_{\sigma_{SP}(t)}$. Note that P may contain new and old predicates.

Algorithm 1 shows how the `FastAbstraction` and `PreciseAbstraction` are combined. It first computes the approximated abstraction (line 4). When a spurious counterexample is encountered as a result of the model checking (line 6), the spurious transitions are removed by using the precise abstraction technique (line 2) with the predicates returned by `SpuriousTransition` (line 9). If no spurious transitions are found, the spurious path is removed by

using the precise abstraction technique (line 20) with the predicates returned by `SpuriousPath` (line 6).

3.1 Refining spurious transitions (lines 9–13)

Suppose some transitions $t_1, \ldots, t_n$ of the counterexample π found by `ModelCheck` are spurious. This means that the function σ_{ST} returned by `SpuriousTransition` maps those transitions to some non-empty set of predicates. Let us define the clustering of predicates Γ as $\{\sigma_{ST}(t_i)\}_{1 \leq i \leq n}$ (i.e., Γ contains the set of predicates $\sigma_{ST}(t_i)$ for every transition in the spurious counterexample). The spurious transition refinement procedure proceeds as follows: For each cluster, $P \in \Gamma$, the refinement algorithm computes $\hat{T}_P$, which is a precise computation of the abstract transition relation projected on the predicates of the cluster. In order to rule out every spurious transition among $t_1, \ldots, t_n$, the refinement algorithm updates the abstract transition relation as follows:

$$\alpha' := \alpha \wedge \bigwedge_{P \in \Gamma} \hat{T}_P \tag{2}$$

Note that, in general, every cluster, P, is a subset of the global set of predicates, Π. This means that each constraint $\hat{T}_P$ is an over-approximation of the precise abstraction computed over Π. Nevertheless, $\hat{T}_P$ is precise with regard to the predicates P, in the sense that it removes all the unrealistic abstract transitions that can be defined by those predicates.

The following theorem states the soundness of this refinement step:

Algorithm 1: A new abstraction-refinement algorithm combining fast and precise abstractions.

```
1  MixCegarLoop(TransitionSystem M, Property F)
2  begin
3      Π = InitialPredicates(F,T);
4      α = FastAbstraction(T,Π);
5      while not TIMEOUT do
6          π = ModelCheck(α,F);
7          if π = ∅ then return CORRECT;
8          else
9              σ_ST = SpuriousTransition(π);
10             if σ_ST ≠ ∅ then
11                 foreach t ∈ π do
12                     C = PreciseAbstraction(T,σ_ST(t));
13                     α = α ∧ C;
14             else
15                 σ_SP = SpuriousPath(π);
16                 if σ_SP = ∅ then return INCORRECT;
17                 else
18                     foreach t ∈ π do
19                         Π = Π ∪ σ_SP(t);
20                         C = PreciseAbstraction(T,σ_SP(t));
21                         α = α ∧ C;
22  end
```

Theorem 1 *For every spurious transition* t_i, $1 \le i \le n$, $t_i \not\models \alpha'$.

Proof Sketch The proof comes directly from the definition of σ_{ST} (it relies therefore on the soundness of a particular `SpuriousTransition` technique): for $1 \le i \le n$, since $t_i \not\models \hat{T}_{\sigma_{ST}(t_i)}$, $t_i \not\models \alpha'$.

In Sect. 2, we discussed that the techniques used to remove spurious transitions require adding a constraint to the abstract transition relation. The Das-Dill technique removes only *one* abstract spurious transition per refinement iteration. When the abstraction is built with a high level of approximation, this technique is highly inefficient because it requires a large number of iterations. The UNSAT core can be used to generate a more relaxed constraint that removes more spurious transitions in one iteration of the CEGAR loop. It can even remove some that are not present in the spurious counterexample. However, it highly depends on the heuristic to cut the UNSAT proof and it is still tightly coupled with the spurious counterexample. By using the precise component $\hat{T}_P$, we remove all spurious transitions which can be expressed with combinations of the predicates in P. This is much stronger than the standard techniques (and, of course, computationally more expensive).

3.2 Refining spurious paths (lines 15–21)

We adopt the cluster-based approach described above to the removal of the spurious path. Our technique uses `SpuriousPath` to produce the set of predicates that are sufficient to rule out the spurious counterexample. The set of predicates generated by the standard predicate-discovery techniques (described in Sect. 2) includes both current predicates and new predicates, that together rule out the spurious counterexample. Our technique considers this set of old and new predicates as a new cluster.

Suppose the path $t_1, \ldots, t_n$ to be spurious. This means that the function σ_{SP} returned by `SpuriousPath` maps each t_i to some non-empty set of predicates. Let us define the clustering of predicates Γ as $\{\sigma_{SP}(t_i)\}_{1 \le i \le n}$ (i.e., Γ contains the set of predicates $\sigma_{SP}(t_i)$ for every transition in the spurious counterexample). The computation of the updated abstract transition relation is identical to spurious transition case, i.e.

$$\alpha' := \alpha \wedge \bigwedge_{P \in \Gamma} \hat{T}_P \tag{3}$$

Note that this time, unlike the case of spurious transitions, the clusters involve new predicates.

By definition, the set of predicates produced by `SpuriousPath` is sufficient to remove the spurious counterexample only if the precise abstraction is used. In fact, spurious transitions over such predicates (possibly created by the approximation abstraction) might create the same

spurious counterexample. Our technique guarantees that this does not happen. This is achieved by using the precise component $\hat{T}_P$.

The following theorem states the soundness of this refinement step:

Theorem 2 *For every spurious path* π, $\pi \not\models \alpha'$.

Proof Sketch The proof comes directly from the definition of σ_{SP} (it relies therefore on the soundness of a particular `SpuriousPath` technique).

In Sect. 2, we referred to the different techniques used to refine the set of predicates. These are orthogonal to the way the abstract transition relation is updated with the new predicates. This is typically done with the same procedure used to compute the initial abstract transition relation given the initial set of predicates. Here, we add a constraint whose precision is determined by the clustering obtained with the spurious path. Thus, it is more precise than `FastAbstraction` but less precise than `PreciseAbstraction`.

3.3 Localized abstraction

The algorithm shown in Algorithm 1 was defined for a monolithic transition relation. When the set of predicates returned by the `SpuriousTransition` or `SpuriousPath` procedures covers the whole set Π of current predicates, the constraint that `MixCegarLoop` adds to the abstract transition corresponds exactly to the precise abstraction. This way, the

---

Algorithm 2: "Synergy" algorithm with localized abstraction.

```
1  MixCegarLoop(TransitionSystem M, Property F)
2  begin
3    foreach T in M do Π(T) = InitialPredicates(F,T);
4    foreach T in M do α(T) = FastAbstraction(T,Π);
5    while not TIMEOUT do
6      π = ModelCheck(α,F);
7      if π = ∅ then return CORRECT;
8      else
9        σ = SpuriousTransition(π);
10       if σ ≠ ∅ then
11         foreach t ∈ π do
12           T = τ(t);
13           C = PreciseAbstraction(T,σ(t));
14           α(T) = α(T) ∧ C;
15       else
16         σ_SP = SpuriousPath(π);
17         if σ_SP = ∅ then return INCORRECT;
18         else
19           foreach t ∈ π do
20             T = τ(t);
21             Π(T) = Π(T) ∪ σ_SP(t);
22             C = PreciseAbstraction(T,σ_SP(t));
23             α(T) = α(T) ∧ C;
24 end
```

abstraction refinement becomes as expensive as `Precise-Abstraction`. We limit this disadvantage by localizing the abstraction to some parts of the program. Some software model checkers (e.g., BLAST [23] and SATABS [15]) use the control flow graph as a partitioning of the transition relation to implement such localization. During the abstraction refinement, they keep a set of predicates and an abstract transition relation for each program location, and perform the abstraction for each local transition relation separately.

Our algorithm implements the localized procedure as part of the CEGAR algorithm as shown in Algorithm 2. The algorithm treats the system M as a set of concrete transition relations, one for every location of the control-flow graph. For each transition relation T, it computes an abstract transition relation $\alpha(T)$ (line 4); when a spurious counterexample is encountered as a result of the model checking (line 6), spurious transitions and path are removed by using the precise abstraction technique (line 13 and 22). The difference from the monolithic case (presented earlier in this section) is that in the localized version, every transition t of the spurious counterexample π is associated with a particular abstract transition relation, denoted $\tau(t)$. Thus, when the refinement step of the algorithm has to add a new constraint, it changes only the transition relation corresponding to either the spurious transition (as part of the spurious transition refinement step, lines 9–14) or to each transition of the spurious path (as part of the spurious path refinement step, lines 16–23).

By exploiting the localized-abstraction framework, the algorithm reduces the abstraction computation to the parts of the system that are relevant to the property and keeps the approximated abstraction in all parts of the program that are irrelevant to prove the property.

4 Evaluation

We implemented the proposed algorithm in the framework of software model checking. We used the SATABS [15] model checker as a platform for our experiments. As described in Sect. 3, the new CEGAR loop uses four subroutines. We experimented with the following techniques implemented in SATABS:

– for `FastAbstraction`, we used a fast abstraction technique based on the computation of the weakest precondition; it assigns to the next predicate its weakest precondition if this is a current predicate; it does not allow a general Boolean combination of predicate variables;
– for `PreciseAbstraction`, we used a precise abstraction based on the enumeration of possible transitions by means of a SAT solver: we force the SAT solver to find all the solutions of the quantifier-elimi-

nation problem by iteratively adding the negation of previous assignments as clauses [15];
– for `SpuriousTransition`, we used the SAT-based technique of [24] [2]; this calls a SAT solver to check if a transition is spurious; if the transition is not realistic, it inspects the UNSAT proof to find the relevant predicates;
– for `SpuriousPath`, we used a technique based on weakest precondition; it computes the weakest preconditions of the current predicates along the transitions of the spurious path; it uses these expressions to produce a set of current and new predicates that are sufficient to rule out the spurious path.

The SAT solver used by `PreciseAbstraction` and `SpuriousTransition` was MiniSAT [19].

We implemented the new algorithm and enhanced SAT-ABS with two new procedures: the first (we will refer to it as `NewST`) affects how the abstraction is refined in the case of spurious transitions, as described in Sect. 3.1; the second (`NewSP`) refines the abstraction in the case of spurious paths, as described in Sect. 3.2.

We compared the new algorithm with the abstraction-refinement loop based on the pure fast abstraction (referred as `WP`) and the pure precise abstraction (referred as `SATQE`) using the standard SATABS implementations of latter techniques. The new algorithm was evaluated with either `New-SP` or `NewST` or both together. Thus, in case `NewSP` was not used, the default refinement of SATABS based on fast abstraction was used.

We ran the experiments on a AMD Dual-Core Opteron 2212 machine with 2 GHz CPU and Ubuntu 7.04. The techniques were evaluated on the sets of ANSI-C programs as benchmarks [3] with different assertions in it. For every experiment, we verified one property at a time.[4]

4.1 Shopping agent benchmark

We first compared the different techniques on a C implementation of a multi server/client shopping agent system (described in details in [7]) as reported in Fig. 2. This example is particularly interesting because the fast abstraction produces a number of spurious transitions exponential in the number of predicates

---

[2] We also experimented with an implementation of technique [17], but it reached 200 CEGAR iterations even on the small examples.

[3] Complete version of results as well as tools and examples are available at http://www.verify.inf.usi.ch/projects/synergy.

[4] We observed that verifying several assertions at the same time may affect the comparison in a unreliable way, since the counterexample produced by the model checker may vary according to different abstract models. This way, at the same iteration we might obtain different predicates which might close the CEGAR loop in a different number of iterations.
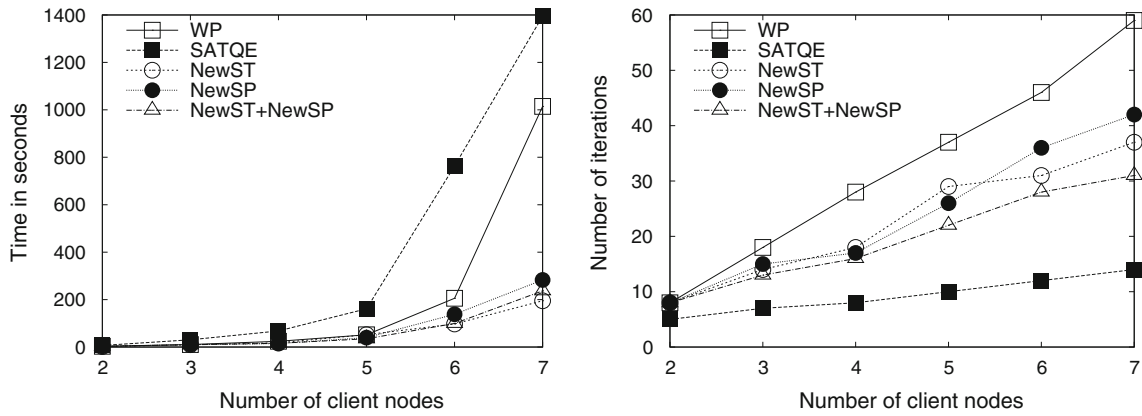
Fig. 2 Total running time in seconds (*left*) and number of iterations (*right*) plotted against the number of clients
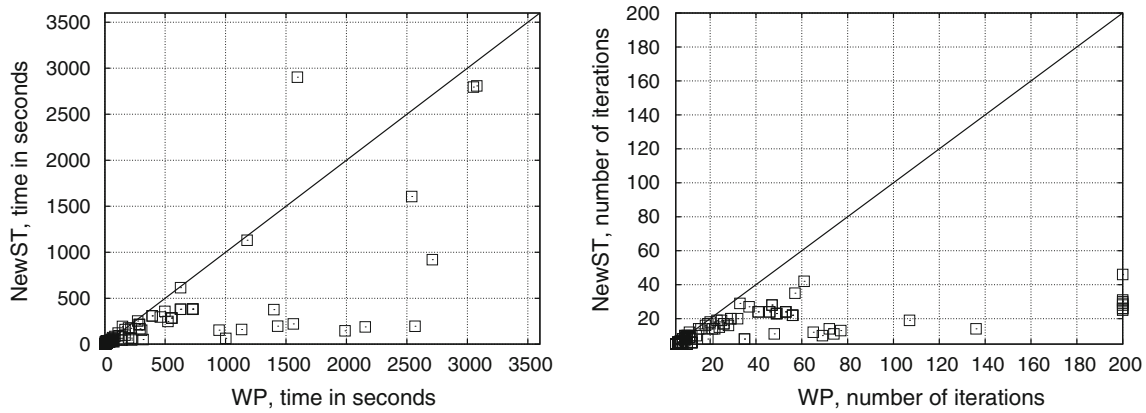

Fig. 3 Comparison of time in seconds (*left*) and number of iterations (*right*) used by WP and NewST

As seen in Fig. 2, the performance of the weakest-precondition-based (WP) and the SAT-based abstractions (SATQE) is comparable. Notably, NewST separately and in combination with NewSP is much more efficient than either WP or SATQE. WP and NewSP are sensitive to a number of spurious transitions and, due to the nature of the example, grow exponentially with the growth of the model. NewST efficiently removes spurious transitions and significantly reduces the number of iterations. In Fig. 2 (right) we note that the new technique as expected has a balanced number of iterations between WP and SATQE. This produces an evident saving in time (as shown in Fig. 2 left) comparing with either WP (up to factor of 5) and with SATQE (up to factor of 7).

4.2 Benchmark test suite from Ku et al.

Next, we evaluated the techniques on the benchmark set proposed in [28]. For this benchmark set the authors collected a large number of large-scale C programs with known buffer-overflow bugs and their fixed versions. The test suite includes applications such as Sendmail, Apache HTTP server, Samba

etc., though, the original programs were stripped down by substituting libraries with stubs. The benchmark set contains 568 [5] test cases, of which 261 are fixed versions of the programs.

4.2.1 Overall results

We limited the execution with 1 hour or 200 iterations of CEGAR per test case. Under this threshold 377 test cases completed by at least one of the techniques. In fact, 40% of them were completed in less than 2 seconds by all techniques and not more than 5 iterations. For this test cases the performance difference was not relevant and we exclude them from the comparison charts (if the opposite is not stated explicitly). For the remaining test cases SATABS needs on average 42 predicates to perform a check, with a maximum of 177 predicates.

Only NewST was able to complete all of 377 considered test cases. WP did 9 less, while SATQE and NewSP failed to

---

[5] We reported to the benchmark authors that 17 test cases are incorrect and 31 test cases do not pass correctly through our front-end; thus, only 520 test cases were used.

finish within a given limit on 76 and 26 test cases, respectively.

4.2.2 WP vs. NewST

The notable comparison of two most effective methods — WP and NewST — gives a better understanding of the advantage of the new techniques. Figure 3 reports the scatter plots of the comparison. The results show that NewST almost systematically outperforms WP. In 98% of the test cases it requires fewer iterations to verify the property. Smaller number of iterations leads to reduction of the total verification time for 53% of the tests. On average, it decreased the total time by 42%, reaching more than double performance gain for some cases. For the small test cases (i.e. 5–10 iterations to complete) the application of the new technique does not give any significant advantage, but it becomes more pronounced with the growth of the test case complexity. The more time the model checking step in CEGAR requires, the bigger reduction in total time the CEGAR loop obtains due to fewer iterations.

4.2.3 Setting up a threshold for PreciseAbstraction

In 47% of the test cases, where NewST was not better than WP, the difference in verification time usually was not bigger than 15%. As an exception, we found only one test case, in which advantage in the smaller number of iterations was not able to compensate for the additional time spent for refinement (the point above the diagonal line in Fig. 3, left).

We investigated the test case: for several program locations PreciseAbstraction computation took longer than the time saved from the reduction in refinement iterations. This was due to the fact that the SAT-based enumeration of all spurious transitions was exponential in the number of predicates returned by SpuriousTransition (or SpuriousPath). Although there were only few transitions where it became critical, we decided to implement a heuristic, which would limit the application of precise computation. The heuristic forbids the application of PreciseAbstraction when the number of predicates reaches a given threshold N_σ. In such cases, FastAbstraction is applied instead of PreciseAbstraction. The value of the threshold depends on the application and the effectiveness of the predicate discovery techniques as well as implementation of PreciseAbstraction and FastAbstraction.

The idea can be further modified to use the already known threshold values. Separate limits can be set for PreciseAbstraction in the SpuriousTransition and SpuriousPath branches. In our experiments we used the pre-computed thresholds that seem optimal for the current implementation of the procedure: we use $N_{\sigma_{ST}} = 13$ for the call of PreciseAbstraction dedicated to the removal

of spurious transition, while $N_{\sigma_{SP}} = 17$ when PreciseAbstraction is used to rule out spurious paths.

We can further optimize this approach by computing the threshold on-the-fly by limiting the maximum execution time for PreciseAbstraction: when the time-out is reached, the number of predicates that made the procedure blow up is used as a new threshold. The approach is shown in Algorithm 3.

We evaluated NewST with the pre-computed thresholds on the test suite from Ku et al. and obtained even better results than for pure NewST. The comparison with WP (Fig. 4) shows that with the heuristic the improvement with NewST is systematic. The comparison between NewST with and without the threshold is shown in Fig. 5. As expected, results

---

Algorithm 3: "Synergy" algorithm with localized abstraction and on-the-fly threshold computation. N_{TO} — time-out value for the PreciseAbstraction; N_σ — computed threshold value; *TimeoutWasReached* — flag, which tracks if PreciseAbstraction was stopped by time-out N_σ.

---

1 MixCegarLoop(*TransitionSystem M, Property F, Time N_{TO}*)
2 **begin**
3   N_σ = *unknown*;
4   **foreach** T *in M* **do** $\Pi(T)$ = InitialPredicates(F,T);
5   **foreach** T *in M* **do** $\alpha(T)$ = FastAbstraction(T,Π);
6   **while** *not TIMEOUT* **do**
7   π = ModelCheck(α,F);
8   **if** $\pi = \emptyset$ **then** return CORRECT;
9   **else**
10   σ_{ST} = SpuriousTransition(π);
11   **if** $\sigma_{ST} \neq \emptyset$ **then**
12   **foreach** $t \in \pi$ **do**
13       $T = \tau(t)$;
14   **if** $N_\sigma = unknown$ *or size*($\sigma_{ST}(t)$) $< N_\sigma$ **then**
15   C = PreciseAbstraction($T,\sigma_{ST}(t),N_{TO}$);
16   **if** *TimeoutWasReached* **then**
17   C = FastAbstraction($T,\sigma_{ST}(t)$);
18         $N_\sigma = size(\sigma_{ST}(t))$;
19   **else**
20   C = FastAbstraction($T,\sigma_{ST}(t)$);
21       $\alpha(T) = \alpha(T) \wedge C$;
22   **else**
23   σ_{SP} = SpuriousPath(π);
24   **if** $\sigma_{SP} = \emptyset$ **then** return INCORRECT;
25   **else**
26   **foreach** $t \in \pi$ **do**
27        $T = \tau(t)$;
28        $\Pi(T) = \Pi(T) \cup \sigma_{SP}(t)$;
29   **if** $N_\sigma = unknown$ *or size*($\sigma_{SP}(t)$) $< N_\sigma$ **then**
30   C = PreciseAbstraction($T,\sigma_{SP}(t),N_{TO}$);
31   **if** *TimeoutWasReached* **then**
32   C = FastAbstraction($T,\sigma_{SP}(t)$);
33          $N_\sigma = size(\sigma_{SP}(t))$;
34   **else**
35   C = FastAbstraction($T,\sigma_{SP}(t)$);
36        $\alpha(T) = \alpha(T) \wedge C$;
37 **end**

**Fig. 4** Scatter plot of time (*left*) and number of iterations (*right*) used by WP and NewST with a threshold


Fig. 5 Comparison of time (*left*) and number of iterations (*right*) used by NewST and NewST with a threshold

of both techniques are similar in more than 90% of the test cases, because the threshold was never reached and Fast-Abstraction was never applied. When the threshold was reached, the results of NewST with N_σ remained very close to the original NewST. But whenever the precise abstraction computation was a bottleneck, the use of the threshold enabled the use of the cheaper fast abstraction consequently resulting in a smaller computation time. The point below the diagonal line in Fig. 5 (left) corresponds to one of the test cases where it happened. As an overall result NewST with a threshold reduced the total verification time by 5% compared with pure NewST.

4.2.4 SATQE, NewSP *and* NewST + NewSP

As expected, SATQE did not perform efficiently whenever a large number of predicates were involved in abstraction. Although on smaller instances ($\leq$30 predicates on average) it showed good results, on large instance it tended to time-out. Thus, it completed 76 test cases less than NewST. NewSP

performed better (only 26 test cases were not finished) but still was worse than WP and NewST. The cause of the problem was similar to the one of SATQE or of NewST without a threshold: NewSP obtained too many predicates from SpuriousPath and the precise computation became very expensive. Nevertheless, it scaled better than SATQE — see Fig. 6 for comparison. Notice that both techniques required fewer iterations than NewST and WP (Fig. 3).

The combination of NewST and NewSP outperformed NewSP (Fig. 7). But the usage of PreciseAbstraction also caused the problem here and did not allow to compete against NewST. Therefore, a threshold for NewSP was also applied similar to its use in the NewST branch (Algorithm 3, lines 2–3).

We compared the fastest technique so far, NewST with a threshold, and a combination of NewSP and NewST with thresholds (Fig. 8). However, on our test suite the winner was not obvious. Although NewST + NewSP variant got more information from counterexamples to remove the spurious behaviors with (likely) cheap computation, the advantage

**Fig. 6** Scatter plot of time (*left*) and number of iterations (*right*) used by `SATQE` and `NewSP`

Fig. 7 Scatter plot of time (*left*) and number of iterations (*right*) used by `NewSP` and `NewST + NewSP`
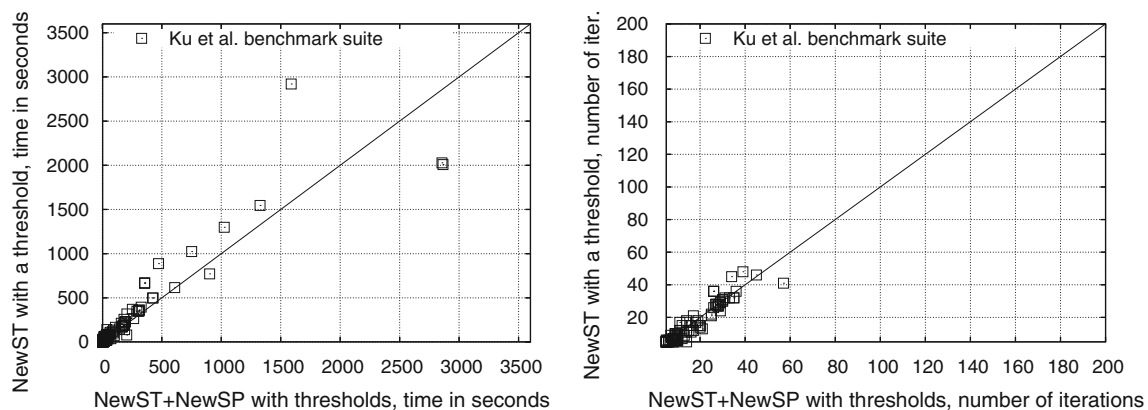

Fig. 8 Scatter plot of time (*left*) and number of iterations (*right*) used by `NewST` and `NewST + NewSP` with thresholds

over `NewST` was not enough to compensate for the additional call to precise abstraction computation. Nevertheless, it confirmed that the use of a threshold helped to avoid problems caused by `PreciseAbstraction`.

4.3 Evaluation on large-scale programs

We experimented with the various large-scale programs from the open-source software packages like `INN`, `WU-FTPD`, `GnuPG` and others. [6] We applied the most effective methods—`WP`, `NewST` and `NewST +NewSP` with thresholds—and analyzed the programs for memory bounds violations.

The overall results on average repeated those from the benchmark suite with an exception that real programs had fewer trivial assertions. Here, we report the outcome for one

---

[6] All the benchmarks were taken from http://www.cprover.org/goto-cc/.

Table 2 SATABS's total time and number of refinement iterations on a set of claims obtained for `inn-encode` `2.4.3` program

#	Total time			Number of iterations		
	WP	NewST	NewST + NewSP	WP	NewST	NewST + NewSP
1	3.478	3.464	2.871	5	5	4
2	2.243	2.318	1.892	4	4	3
3	7.977	8.345	6.64	6	6	5
4	124.013	104.657	83.893	25	19	10
5	4.149	4.222	3.529	4	4	3
6	137.317	97.449	121.919	28	17	12
7	2.683	2.698	1.567	3	3	2
8	2.712	2.636	1.594	3	3	2
9	37.86	28.783	31.429	10	8	7
10	27.575	27.225	29.612	9	9	7
11	5.975	5.801	4.727	6	6	4
12	76.945	49.822	71.106	13	10	10
13	TO	TO	TO	TO	TO	TO
14	7.894	8.195	6.985	6	6	5
15	128.271	98.01	88.266	26	19	10
16	4.207	4.261	3.42	4	4	3
17	145.884	112.898	122.006	30	19	13
18	2.113	2.123	1.33	3	3	2
19	2.193	2.158	1.37	3	3	2
20	31.598	22.788	27.131	9	7	6
21	27.163	22.906	28.05	10	8	6
22	4.349	4.495	3.111	5	5	3
23	77.919	49.293	67.942	13	10	10
24	10.981	9.494	11.103	8	7	6
25	7.408	7.603	6.62	6	6	5
26	150.855	123.884	112.992	32	23	14
27	4.439	4.393	3.592	4	4	3
28	125.827	73.21	97.236	30	15	14

TO stands for time-out (3,600 s)

of the experiments. We analyzed the `encode` program from the `inn` utilities suite version 2.4.3 [1]. It produces a seven-bit printable encoding of `stdin` on `stdout` and serves as a good example of a small memory-operating piece of C code. This program was taken as an example also because it is not very big (1.1KLOC) and has only 28 locations where a safety of the memory access should be checked. The size of the program allowed most of the claims to be verified within one hour time limit.

The results are reported in Table 2. For each claim and each technique we showed a total verification time and a number of the required refinement iterations. As expected, the reduction in the refinement iterations resulted in reduction of the total verification time. `NewST` used fewer refinements than `WP` in 12 out of 28 claims and won in verification time as well. Interesting to notice, the advantage was achieved any time more than 10 refinement iterations

were required. For other 16 claims two techniques showed approximately the same result. Precise abstraction computation was localized and never required a significant time. `NewST` + `NewSP` required fewer refinements than `WP` in all 28 claims and, as a result, it outperformed `WP` on all but 3 claims. However, it did not perform better than `NewST` on every claim and therefore they are comparable in their advantages.

5 Conclusions and future work

We presented a new approach to the abstraction refinement that combines precise and approximated techniques. On the one hand, the proposed algorithm benefits from the precise component, because it avoids too many iterations due to spurious transitions of the abstract model. On

the other hand, it uses the fast component to discover the spurious counterexample. Moreover, by exploiting the localized-abstraction framework, it reduces the abstraction computation to the parts of the system that are relevant to the property and keeps the approximated abstraction in all parts of the program that are irrelevant to prove the property. Our technique is independent of any particular abstraction or refinement procedure and can be used for any combination of the existing abstraction and refinement techniques.

We performed an extensive evaluation on large-scale programs comparing the new technique with the classical precise and imprecise algorithms. Our tests with various benchmarks show that the new approach systematically outperforms both precise and imprecise techniques. Altogether, it confirms that our new technique achieves the goal of reducing the number of iterations of the CEGAR loop.

In this paper, the goal of the experimental evaluation was to validate the new technique on spurious transition refinement. Thus, we maintained the same tool framework and we did not change orthogonal techniques such as predicate discovery. As a future work, we are interested in implementing the same approach in other tools such as BLAST [23] and in integrating it with interpolation-based approaches to predicate discovery [22,27]. Another interesting direction is to investigate the same trade-off between precise and approximated approaches in the context of purely interpolation-based model checking [31] which does not need predicate abstraction. Also we plan to establish fine-grained correspondence between the semantics of the analyzed model (e.g. semantic of C code instructions) and the combination of fast/precise abstraction.

References

1. https://www.isc.org/software/inn
2. Ball, T., Cook, B., Das, S., Rajamani, S.K.: Refining approximations in software predicate abstraction. In: TACAS 388–403 (2004)
3. Ball, T., Majumdar, R., Millstein, T.D., Rajamani, S.K.: Automatic predicate abstraction of C programs. In: PLDI 203–213 (2001)
4. Ball, T., Podelski, A., Rajamani, S.K.: Boolean and Cartesian abstraction for model checking C programs. STTT **5**(1), 49–58 (2003)
5. Ball, T., Rajamani, S.K.: Boolean programs: A model and process for software analysis. Technical report 2000–2014, Microsoft research, February (2000)
6. Ball, T., Rajamani, S.K.: Generating abstract explanations of spurious counterexamples in C programs. Technical report 2002–2009, Microsoft research, September (2002)
7. Braghin, C., Sharygina, N., Barone-Adesi, K.: Automated verification of security policies in mobile code. In: Davies, J., Gibbons, J., (eds) IFM. volume 4591 of Lecture Notes in Computer Science. Springer, Berlin, pp. 37–53 (2007)
8. Bryant, R. E.: Graph-based algorithms for boolean function manipulation. IEEE Trans Comput **C-35**(8), 677–691 (1986)
9. Burch, J.R., Clarke, E.M., McMillan, K.L., Dill, D.L., Hwang, L.J.: Symbolic model checking: 10^{20} states and beyond. Inf. Comput. **98**(2), 142–170 (1992)
10. Cavada, R., Cimatti, A., Franzén, A., Kalyanasundaram, K., Roveri, M., Shyamasundar, R. K.: Computing predicate abstractions by integrating BDDs and SMT solvers. In: FMCAD IEEE, pp. 69–76 (2007)
11. Clarke, E., Talupur, M., Veith, H., Wang, D.: SAT based predicate abstraction for hardware verification. In: SAT (2003)
12. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-Guided Abstraction Refinement. In: CAV, pp. 154–169 (2000)
13. Clarke, E.M., Grumberg, O., Long, D.E.: Model checking and abstraction. ACM Trans. Program. Lang. Syst. **16**(5), 1512–1542 (1994)
14. Clarke, E.M., Gupta, A., Kukula, J.H., Strichman, O.: SAT based abstraction-refinement using ILP and machine learning techniques. In: CAV, pp. 265–279 (2002)
15. Clarke, E.M., Kroening, D., Sharygina, N., Yorav, K.: Predicate abstraction of ANSI-C programs using SAT. Formal Methods in System Design **25**(2–3), 105–127 (2004)
16. Colón, M., Uribe, T.E.: Generating finite-state abstractions of reactive systems using decision procedures. In: CAV, pp. 293–304 (1998)
17. Das, S., Dill, D.L.: Successive approximation of abstract transition relations. In: LICS, pp. 51–60 (2001)
18. Das, S., Dill, D.L., Park, S.: Experience with predicate abstraction. In: CAV (1999)
19. Eén, N., Sörensson, N., An extensible sat-solver. In: SAT, pp. 502–518 (2003)
20. Graf, S., Saïdi, H.: Construction of abstract state graphs with PVS. In: CAV, pp. 72–83 (1997)
21. Gupta, A., Strichman, O.: Abstraction refinement for bounded model checking. In: CAV, pp. 112–124 (2005)
22. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. In: POPL, pp. 232–244 (2004)
23. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy Abstraction. In POPL, pp. 58–70 (2002)
24. Jain, H., Kroening, D., Sharygina, N., Clarke, E.M.: Word level predicate abstraction and refinement for verifying RTL verilog. In: DAC, pp. 445–450 (2005)
25. Jain, H., Ivancic, F., Gupta, A., Ganai, M. K.: Localization and register sharing for predicate abstraction. In: TACAS, pp. 397–412 (2005)
26. Jhala, R., McMillan, K.L.: Interpolant-based transition relation approximation. In: CAV, pp. 39–51 (2005)
27. Jhala, R., McMillan, K.L.: A practical and complete approach to predicate refinement. In: TACAS, pp. 459–473 (2006)
28. Ku, K., Hart, T. E., Chechik, M., Lie, D.: A buffer overflow benchmark for software model checkers. In: ASE '07 ACM Press, pp. 389–392 (2007)
29. Lahiri, S.K., Ball, T., Cook, B.: Predicate abstraction via symbolic decision procedures. Log. Methods Comput. Sci. **3**(2) (2007)
30. Lahiri, S.K., Nieuwenhuis, R., Oliveras, A.: SMT techniques for fast predicate abstraction. In: CAV. LNCS, Springer, Berlin, pp. 424–437 (2006)
31. McMillan, K.L.: Lazy abstraction with interpolants. In: CAV, pp. 123–136 (2006)
32. McMillan, K.L.: Applying SAT methods in unbounded symbolic model checking. In: CAV, pp. 250–264 (2002)
33. Nielson, F., Nielson, H. R., Hankin, C. L.: Principles of Program Analysis. Springer, Berlin (1999)
34. Sharygina, N., Tonetta, S., Tsitovich, A.: The synergy of precise and fast abstractions for program verification. In: 24th annual ACM symposium on applied computing. Honolulu, Hawaii, USA, ACM (2009)