

# Counterexample-Guided Precondition Inference<sup>\*</sup>

Mohamed Nassim Seghir and Daniel Kroening

Computer Science Department, University of Oxford

**Abstract.** The precondition for an assertion inside a procedure is useful for understanding, verifying and debugging programs. As the procedure might be used in multiple calling-contexts within a program, the precondition should be sufficiently general to enable re-use. We present an extension of counterexample-guided abstraction refinement (CEGAR) for automated precondition inference. Starting with an over-approximation of both the set of safe and unsafe states, we iteratively refine them until they become disjoint. The resulting precondition is then necessary and sufficient for the validity of the assertion, which prevents false alarms. We have implemented our approach in a tool called P-Gen. We present experimental results on string and array-manipulating programs.

## 1 Introduction

Software model checking is a popular technique for program verification. A diverse range of tools based on this approach have been developed (e.g., SLAM [1], BLAST [20], MAGIC [7], SATABS [8] and TERMINATOR [11]) and successfully applied to real-world software. The key to effectiveness of these tools is *abstraction*, and predicate abstraction [16] is a well-established instance. The predicate discovery in tools implementing it is driven by *counterexample-guided abstraction refinement* [9], commonly known as CEGAR.

Most of the tools above answer the usual verification question: “given an assertion at some program location, is this assertion always valid?” When considering just a fragment of a program containing an assertion, we can ask a slightly different question: “In which context is the assertion valid?” The code fragment might be a procedure that is called at different program locations, hence the computed context should be as general as possible to be reusable at the different call sites. A simple and straightforward way to infer a precondition is to compute a conservative abstraction of the set of unsafe states, i.e., those states that can reach an error, and using its complement as precondition. The problem with this approach is that an over-approximation of the set of unsafe states might include safe states as well, resulting in an over-conservative precondition. The abstraction must then be refined by removing some of the safe states. This cannot always be performed in an enumerative fashion, as the set of safe states is often infinite.

---

<sup>\*</sup> Supported by ERC project 280053, EPSRC project EP/H017585/1, the EU FP7 STREP PINCETTE and the ARTEMIS VETESS project.

We propose a solution to this problem based on the abstraction (and thus generalization) of both the set of safe and unsafe states. Our approach is based on the CEGAR paradigm: starting with an over-approximation of both sets, we iteratively refine them until they become disjoint. Thus, the resulting precondition is sufficient and also necessary for the validity of the assertion. This guarantees the absence of false alarms, as the violation of the precondition by some calling context entails the violation of the assertion within the procedure. Our contributions are summarized as follows:

- A novel approach to generate *exact* preconditions, i.e., necessary and sufficient. Thus, the precondition is independent from the calling context. Most of the approaches in the literature generate preconditions that are only sufficient, thus the precondition often has to be re-adjusted if it is not satisfied by some calling context. In our case, a violation of the precondition will result in a real error, and we thus avoid false alarms.
- An implementation of the approach using ingredients that are common to most CEGAR-based verification tools. Thus, our technique represents a generic scheme for extending other tools to infer preconditions.
- A simple predicate inference mechanism for algorithms that manipulate arrays used on top of the standard predicate refinement procedure. This simple technique generates predicates that are often adequate to obtain the right program invariant and subsequently obtain the desired precondition.

The remainder of this paper is organized as follows: Section 2 illustrates our approach by means of examples. Section 3 introduces background material. Section 4 describes our approach for precondition inference and the refinement technique used in the CEGAR loop. Section 5 presents experimental results and Section 6 discusses related work.

## 2 Examples

Consider the program `copy` given in Figure 1(a). It takes as parameters two arrays  $a$  and  $b$  and the length  $bJ$  of array  $b$ . The program copies the elements of array  $b$  in the range  $\{0, \dots, bJ - 1\}$  to the corresponding range in array  $a$ . The access to array  $a$  is safe if the index expression is in the range  $\{0, \dots, aJ - 1\}$ , where  $aJ$  is the length of array  $a$ . It is trivial to see that the lower bound is not violated. Let us then focus on the upper bound. The safety condition with regards to the upper bound is expressed by the assertion at location  $\ell_2$ . Our goal is to find a precondition for procedure `copy` that guarantees that this assertion is never violated. The precondition must be expressed only using program elements visible at the entry-point of the procedure, i.e., it must be a predicate over the procedure parameters and the global variables. The precondition should also be *exact*, i.e., it should neither be too strong nor too weak.

*Transformation to reachability.* We will now illustrate our approach to precondition inference. We use standard notation and formally represent programs in

<pre> <b>void</b> copy(<b>int</b> a[], <b>int</b> b[], <b>int</b> b_l) {     <b>int</b> i;     ℓ<sub>0</sub> : i = 0;     ℓ<sub>1</sub> : <b>while</b>(i &lt; b_l)     {         ℓ<sub>2</sub> : <b>assert</b>(i &lt; a_l);             a[i] = b[i];             i++;     } } </pre> <p style="text-align: center;">(a)</p>	<pre> <b>void</b> copy_2(<b>int</b> a[], <b>int</b> b[]) {     <b>int</b> i;     ℓ<sub>0</sub> : i = 0;     ℓ<sub>1</sub> : <b>while</b>(b[i] != 0)     {         ℓ<sub>2</sub> : <b>assert</b>(i &lt; a_l);             a[i] = b[i];             i++;     } } </pre> <p style="text-align: center;">(b)</p>
---	--

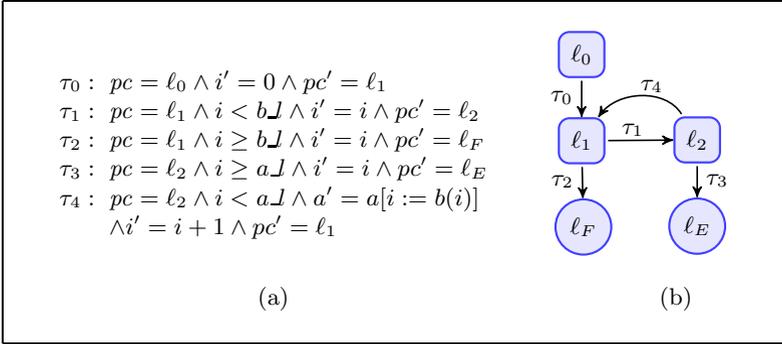
**Fig. 1.** Two simple programs that copy a range of elements from array  $b$  to array  $a$ . In procedure `copy`, the limit of the range to be copied is explicitly given via  $b\_l$ . In `copy_2`, the range is implicitly delimited via the sentinel value 0.

terms of transition constraints over primed and unprimed program variables. The set of transition constraints corresponding to program `copy` (Figure 1(a)) is given in Figure 2(a) and the associated control flow graph is given in Figure 2(b). The program counter is modeled explicitly using the variable  $pc$ , which ranges over the set of control locations. The assertion in the original program is replaced with a conditional branch whose condition is the negation of the assertion and whose target is the error location  $\ell_E$ . The special location  $\ell_F$  is the *final location*, and has no successor.

Observe that the error location is only reachable if  $i \geq a\_l$  evaluates to true at location  $\ell_2$ . The final location  $\ell_F$  is reached in paths without error. The transition  $\tau_0$  corresponds to the initialization of variable  $i$ . The transition  $\tau_1$  represents the entrance to the loop and  $\tau_2$  the exit from the loop. The assertion is modeled via the transition  $\tau_3$ , which conditionally alters the control flow to the error location. Finally, the transition  $\tau_4$  models the remainder of the loop body after the `assert` statement. Arrays  $a$  and  $b$  are represented by uninterpreted function symbols, and  $a[x := e]$  denotes function update (the expression is equal to  $a$  where the  $x^{\text{th}}$  element has been replaced by  $e$ ).

*Over-approximating the unsafe states.* It is in general not possible to enumerate all the traces of a program. In our example, the program contains a cycle  $\langle \tau_1; \tau_4 \rangle$  (Figure 2(b)) that can be unfolded an indefinite number of times, leading to an infinite number of traces. A solution to this problem is to provide a *backwards inductive invariant*: an invariant that includes all error states and which is inductive under the application of  $\text{wp}^1$ . Predicate abstraction [16] is a suitable

<sup>1</sup>  $\text{wp}(\tau, \varphi)$  is the weakest precondition for the formula  $\varphi$  with respect to statement (transition constraint)  $\tau$ . It extends to a sequence of statements (trace)  $\pi$ .



**Fig. 2.** Transition constraints for program copy (a) and the corresponding graphical representation (b) or the control flow graph

technique for building such an invariant. The key challenge when applying predicate abstraction is the choice of predicates.

Naive approaches for inferring predicates, e.g., based on weakest preconditions, often diverge [22]. In our example, suppose that we first obtain the error path  $\tau_0; \tau_1; \tau_3$  from the abstract model. The analysis of this path via  $\text{wp}$ , i.e.,  $\text{wp}(\tau_0; \tau_1; \tau_3, pc = \ell_E)$ , gives us the formula  $0 < b \perp \wedge 0 \geq a \perp$ . We add the predicates  $0 < b \perp$  and  $0 \geq a \perp$  and set the precondition  $\varphi$  to be  $0 \geq b \perp \vee 0 < a \perp$ . If we unfold the loop once more we obtain the error trace  $\tau_0; \tau_1; \tau_4; \tau_1; \tau_3$ . We have  $\text{wp}(\tau_0; \tau_1; \tau_4; \tau_1; \tau_3, pc = \ell_E) \equiv 1 < b \perp \wedge 1 \geq a \perp$ . This new trace is not covered by the previous one, as it is still feasible under our precondition  $\varphi$ . We then update  $\varphi$  to rule out the new trace to obtain  $(0 \geq b \perp \vee 0 < a \perp) \wedge (1 \geq b \perp \vee 1 < a \perp)$ . After unfolding the loop  $j$  times, we obtain  $\text{wp}(\tau_0; \langle \tau_1; \tau_4 \rangle^j; \tau_1; \tau_3, pc = \ell_E) \equiv j < b \perp \wedge j \geq a \perp$  and the precondition

$$\varphi \equiv \bigwedge_{j>0} j \geq b \perp \vee j < a \perp .$$

We can continue to unfold the loop, every time generating a new trace that is not covered by the previous ones. To address this divergence, we go beyond the syntactic approach to predicate discovery and use techniques to infer more general facts. For example, by linearly combining the predicates  $0 < b \perp$  and  $0 \geq a \perp$  (from the first iteration) we deduce  $a \perp < b \perp$ . This new predicate is a backwards invariant at location  $\ell_0$  with respect to the program, i.e.,

$$\left( \bigvee_{j>0} \text{wp}(\pi_j, pc = \ell_E) \right) \Rightarrow (pc = \ell_0 \Rightarrow a \perp < b \perp) .$$

Thus, the predicate  $a \perp < b \perp$  over-approximates the set of states that reach the error location. The precondition  $\varphi$  is then simply chosen to be  $a \perp \geq b \perp$ , i.e., the complement of that set.

*Over-approximating the Safe States.* When over-approximating the set of states that reach the error location, we always include all error states but may include some safe entry states as well. It means that the precondition, which is the complement of the computed set, may exclude safe traces and is thus unnecessarily strong. To tune the precision of the abstracted set of error states, our new algorithm also over-approximates the set of entry states that reach the final location  $\ell_F$  (the safe states). We then check the intersection of this set with the (over-approximation of the) states that reach the error location. If the intersection is empty, we conclude that our current set of unsafe states does not include any safe state.

As we did for the error location, we obtain the over-approximation of the set of states reaching the final location given as the state formula  $a\perp \geq b\perp \vee 0 \geq b\perp$ . The intersection of this set with the set of unsafe states ( $a\perp < b\perp$ ) is obtained by forming the conjunction:

$$(a\perp \geq b\perp \vee 0 \geq b\perp) \wedge (a\perp < b\perp) .$$

The formula above has satisfying assignments, which means that the two sets are not disjoint. As  $a\perp \geq b\perp$  and  $a\perp < b\perp$  are inconsistent, the intersection can only be in  $0 \geq b\perp$ . Thus, from  $0 \geq b\perp$  we can reach both the error and final location. This outcome is caused by insufficient precision of the abstraction. Let us consider two traces,  $\pi_E = \tau_0; \tau_1; \tau_3$  leading to the error location, and  $\pi_F = \tau_0; \tau_2$  leading to the final location. We have  $\text{wp}(\pi_E, pc = \ell_E) \equiv 0 < b\perp \wedge 0 \geq a\perp$  and  $\text{wp}(\pi_F, pc = \ell_F) \equiv 0 \geq b\perp$ . Thus,  $\pi_E$  is not feasible from states with  $0 \geq b\perp$ , which means that the set of unsafe states is not precise enough. It is then refined by adding the predicate  $b\perp > 0$ , which makes the two sets disjoint. The final precondition  $\varphi$  is given by

$$0 \geq b\perp \vee a\perp \geq b\perp .$$

The precondition  $\varphi$  is now necessary and sufficient, meaning that it does not allow any state to reach the error location and does not exclude any state that reaches the final location.

*Inferring quantified preconditions.* Let us consider a slightly modified version of the previous program `copy`, which is given in Figure 1(b). In this example, the range of elements to be copied from array  $b$  to  $a$  is not explicit, as it is indicated via a sentinel value (0 in the example). After going through the different steps described for the previous example, our method succeeds in inferring the precondition

$$(b[0] = 0) \vee (\exists x \in \{0, \dots, a\perp\}. b[x] = 0) .$$

Observe that the precondition inferred by the algorithm is not equivalent to  $\exists x \in \{0, \dots, a\perp\}. b[x] = 0$ ; it is weaker. This is due the possibility of skipping the loop when  $b[i] \neq 0$  is false, regardless of the value of  $a\perp$ . This implies that runs from states in which  $b[0] = 0$  are safe. This case is expressed via the first disjunct of the precondition above. We will re-visit the second example in Section 4 to illustrate our refinement procedure.

### 3 Preliminaries

In this section, we provide background on counterexample-guided abstraction refinement for predicate abstraction.

*Program.* To aid the formal presentation, we assume that a program is given as a set  $\mathcal{TC}$  of transition constraints  $\tau$ . A transition constraint  $\tau$  is a formula of the form

$$g(X) \wedge x'_1 = e_1(X) \wedge \dots \wedge x'_n = e_n(X) \quad (1)$$

where  $X = \langle x_1, \dots, x_n \rangle$  is a tuple (vector) of program variables, which include the program counter  $pc$ . In (1), unprimed variables refer to the program state before performing the transition and primed ones represent the program state after performing the transition. Formula  $g(X)$  is called the *guard* and the remaining conjuncts of  $\tau$  are the *update* or *assignment*.

*Representing States Symbolically.* Let us write  $\mathcal{V} = \{x_1, \dots, x_n\}$  for the set of variables of the program (including the program counter  $pc$ ). For a variable  $x \in \mathcal{V}$ ,  $Type(x)$  is the type (range) of  $x$  and  $\sigma(x)$  is a valuation of  $x$  such that  $\sigma(x) \in Type(x)$ . The variable  $pc$  ranges over the set of all program locations. Given  $X$  (a tuple of the variables), a program state is the valuation  $\sigma(X) = \langle \sigma(x_1), \dots, \sigma(x_n) \rangle$ .

A set of program states  $S$  is represented symbolically by means of the *characteristic function* of  $S$ . The formula  $\varphi$  represents the set of all those states that correspond to a satisfying assignment of  $\varphi$ , i.e.,  $\{\sigma(X) \mid \varphi[\sigma(X)/X]\}^2$ . We will use sets and their characteristic functions interchangeably. Symbolic states (formulas) are partially ordered via the implication operator  $\Rightarrow$ , i.e.,  $\varphi' \leq \varphi$  means  $\varphi' \Rightarrow \varphi$ .

*State transformer.* For a formula  $\varphi$ , the application of the operator  $\mathbf{pre}$  with respect to the transition constraint  $\tau$  returns a formula representing the set of all predecessor states of  $\varphi$  under the transition constraint  $\tau$ , formally

$$\mathbf{pre}(\tau, \varphi(X)) \equiv g(X) \wedge \varphi[\langle e_1(X), \dots, e_n(X) \rangle / X].$$

For the whole program  $\mathcal{TC}$ ,  $\mathbf{pre}$  is given by

$$\mathbf{pre}(\varphi(X)) \equiv \bigvee_{\tau \in \mathcal{TC}} \mathbf{pre}(\tau, \varphi(X)).$$

For a trace  $\pi = \tau_1; \dots; \tau_n$ , we have

$$\mathbf{pre}(\tau_1; \dots; \tau_n, \varphi) = \mathbf{pre}(\tau_1, \dots, \mathbf{pre}(\tau_{n-1}, \mathbf{pre}(\tau_n, \varphi))).$$

If  $\mathbf{pre}(\pi, \varphi)$  is not equivalent to false, then the trace  $\pi$  is *feasible*.

---

<sup>2</sup> The notation  $f[Y/X]$  represents the expression obtained by replacing all occurrences of every variable from the vector  $X$  in  $f$  with the corresponding variable (value) from  $Y$ . It naturally extends to a collection (set or list) of expressions.

(Un)Safe states To ease the presentation, let us assume that the program contains a single error location  $\ell_E$  and a single final location  $\ell_F$  ( $\ell_E \neq \ell_F$ ).<sup>3</sup> We denote by **bad** the set of *error states*, which is simply given by  $pc = \ell_E$ . Similarly, we call **final** the set of *final states*, which is represented by  $pc = \ell_F$ .

The set of safe states **safe** contains all states from which a final state is reachable. Formally,

$$\mathbf{safe} \equiv \text{lfp}(\mathbf{pre}, \mathbf{final}) \quad (2)$$

where  $\text{lfp}(\mathbf{pre}, \varphi)$  denotes the least fixpoint of the operator **pre** above  $\varphi$ . Similarly, **unsafe** is the set of all states from which an error (**bad**) state is reachable:

$$\mathbf{unsafe} \equiv \text{lfp}(\mathbf{pre}, \mathbf{bad}) . \quad (3)$$

The least fixpoints represent inductive backwards invariants, which we denote by  $\psi_{\mathbf{bad}}$  and  $\psi_{\mathbf{final}}$ , respectively. The invariants are inductive under **pre**, i.e.,

- $\mathbf{bad} \leq \psi_{\mathbf{bad}}$  and  $\mathbf{final} \leq \psi_{\mathbf{final}}$
- $\mathbf{pre}(\psi_{\mathbf{bad}}) \leq \psi_{\mathbf{bad}}$  and  $\mathbf{pre}(\psi_{\mathbf{final}}) \leq \psi_{\mathbf{final}}$

In the absence of non-determinism in the program, the sets of **unsafe** and **safe** states are disjoint, and we have

$$\mathbf{unsafe} \wedge \mathbf{safe} \equiv \mathbf{false} .$$

*Predicate Abstraction.* Predicate abstraction consists of approximating a state  $\varphi$  with a formula  $\varphi'$  constructed as a Boolean combination of predicates taken from a set  $P$ . Here, the term approximation means that any model that satisfies  $\varphi$  must satisfy  $\varphi'$ . Thus, a suitable approximation is obtained via the logical implication “ $\Rightarrow$ ”, i.e.,  $\varphi'$  is the strongest Boolean combination built up from predicates taken from the finite set  $P$  such that  $\varphi \Rightarrow \varphi'$ .

Defining the abstraction function  $\alpha$  as being the strongest Boolean combination of predicates in  $P$  is not practical because of the exponential complexity of the problem. Therefore, we use a lightweight version of  $\alpha$  that consists of building the strongest conjunction of predicates in  $P$ :

$$\alpha(\varphi) \equiv \bigwedge p \mid p \in P \wedge \varphi \Rightarrow p .$$

Let us have  $\mathcal{D}^\sharp$  the domain of formulas built up from the finite set of predicates  $P$ . The domain  $\mathcal{D}^\sharp$  is not closed under **pre**, therefore, we define  $\mathbf{pre}^\sharp$  under which  $\mathcal{D}^\sharp$  is closed. Let us associate the concretization function  $\gamma : \mathcal{D}^\sharp \rightarrow \mathcal{D}$  to  $\alpha$ , we simply choose  $\gamma$  to be the identity function. Functions  $\alpha$  and  $\gamma$  form a *Galois connection* with respect to  $\geq$  ( $\Leftarrow$ ) being the partial order relation for both  $\mathcal{D}$  and  $\mathcal{D}^\sharp$ . Formally speaking

$$\forall x \in \mathcal{D} \forall y \in \mathcal{D}^\sharp . \alpha(x) \geq y \Leftrightarrow x \geq \gamma(y) .$$

---

<sup>3</sup> In case of multiple assertions, we add an edge from each assertion (guarded with the negation of the assertion) to  $\ell_E$ . Similar treatment can be applied in the case of multiple return locations.

Hence, we define  $\text{pre}^\sharp : \mathcal{D}^\sharp \rightarrow \mathcal{D}^\sharp$ , the abstract version of  $\text{pre}$ , as follows:

$$\text{pre}^\sharp(\varphi) \equiv \alpha(\text{pre}(\gamma(\varphi))),$$

and thus

$$\text{pre}^\sharp(\tau, \varphi) = \alpha(\text{pre}(\tau, \varphi)) = \bigwedge p \mid p \in P \wedge \text{pre}(\tau, \varphi) \Rightarrow p.$$

As seen for  $\text{pre}$ , the operator  $\text{pre}^\sharp$  also extends to traces. Henceforth, whenever we write  $\text{pre}_P^\sharp$  we mean that the abstraction (image) is computed by considering predicates from the set  $P$ .

The lattice of abstract states  $(\mathcal{L}, \Rightarrow)$  is finite as the set of predicates is finite. Therefore,  $\text{lfp}(\text{pre}^\sharp, \text{bad})$  ( $\text{lfp}(\text{pre}^\sharp, \text{final})$ ), the least fixpoint for  $\text{pre}^\sharp$  above  $\text{bad}$  ( $\text{final}$ ) in  $\mathcal{L}$ , is computable.

## 4 Precondition Inference

The precondition inference problem can be described as the computation of a formula  $\varphi$  such that:

$$\text{lfp}(\text{pre}, \text{bad}) \wedge \varphi \equiv \text{false} \tag{4}$$

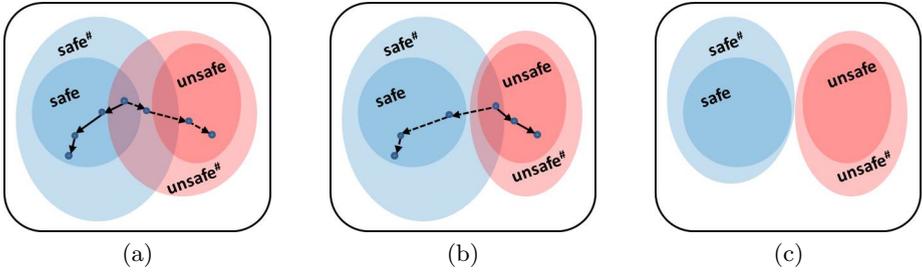
The fixpoint for the preimage-operator  $\text{pre}$  is in general not computable, we thus compute the least fixpoint for  $\text{pre}^\sharp$ . As we have  $\text{lfp}(\text{pre}, \text{bad}) \leq \text{lfp}(\text{pre}^\sharp, \text{bad})$ , it is sufficient to show that

$$\text{lfp}(\text{pre}^\sharp, \text{bad}) \wedge \varphi \equiv \text{false}$$

to conclude the validity of (4). The precondition  $\varphi$  can then be simply chosen as the negation of  $\text{lfp}(\text{pre}^\sharp, \text{bad})$  projected on the entry location. One problem with this approach is that due to the abstraction we may exclude some of the safe, terminating runs. A second challenge is the choice of predicates. We have seen in the illustrative example that a bad choice of predicates can lead to divergence. In what follows, we present a new CEGAR-based algorithm for precondition inference that guarantees that all safe executions are included in the precondition. We also propose a predicate discovery mechanism that goes beyond the approach based on weakest precondition.

### 4.1 Counterexample-Guided Precondition Inference

Our goal is to increase the precision of the set of unsafe states  $\text{unsafe}^\sharp$ , making it free of safe states. This is non-trivial, since we cannot enumerate safe states, as there are in general infinitely many. Hence, we need to construct the set of safe states by over-approximating them as well. Our idea consists of building abstractions of increasing precision of both the set of safe and unsafe states until they become disjoint. We propose an implementation of this idea by extending the classical CEGAR paradigm, where its main ingredients are instantiated in our setting with the following:



**Fig. 3.** Illustration of the main phases of algorithm `InferPrecond`. Dashed arrows indicate that the trace is spurious.

- **Abstraction:** we abstract both the set of safe and unsafe states.
- **Counterexample:** in our context, a counterexample is *two* abstract traces, a safe one and an unsafe one, beginning with a common initial state.
- **Counterexample simulation:** checks if the two traces can be concretized to effectively share a common concrete initial state. This is only possible in the presence of non-determinism in the program. The check is carried out by computing the weakest precondition for each trace. Hence, the counterexample is spurious if the two preconditions are disjoint.
- **Refinement:** the spurious counterexample is ruled out by adding predicates that refine the abstraction such that the two traces cannot share their initial state.
- **Termination criterion:** the iterative process stops when the two abstractions (of safe and unsafe states) are disjoint.

We present algorithm `InferPrecond` (Algorithm 1), which implements a counterexample-guided abstraction refinement loop for precondition inference. The algorithm starts with an over-approximation of both the set of safe and unsafe states (lines 5 and 6), denoted by  $\text{safe}^\#$  and  $\text{unsafe}^\#$ , respectively. It iteratively refines them until their projections onto the initial location become disjoint, i.e.,  $(\text{safe}^\# \wedge \text{unsafe}^\# \wedge pc = \ell_0) \equiv \text{false}$  (Figure 3(c)). The computed precondition is then the set of safe states projected onto the initial location  $\ell_0$  (line 8 of the algorithm).

The refinement process is applied whenever  $\text{safe}^\#$  and  $\text{unsafe}^\#$  intersect, i.e., when we have a bad trace and a safe one sharing their initial state. In Figure 3(a),  $\text{safe}^\#$  and  $\text{unsafe}^\#$  intersect, but the analysis reveals that the initial state is in reality in  $\text{safe}$ , thus the (dashed) trace in  $\text{unsafe}^\#$  is the one that is spurious. After refining  $\text{unsafe}^\#$ , we obtain the abstraction in Figure 3(b). The two sets still intersect, however this time the spurious trace is in  $\text{safe}^\#$ , as the initial state belongs to  $\text{unsafe}$ . The refinement process is carried out by calling the procedure `Refine` at line 14. This procedure takes as parameters two traces, one leading to the error location and another one leading to the final location, and returns a new set of predicates. We describe this procedure in detail in the next section.

**Algorithm 1.** InferPrecond

---

**Input:** set of transition constraints (program)  $\mathcal{TC}$   
**Output:** formula (precondition)

```

1 Var  $P$ : set of predicates;
2 Var  $\text{safe}^\sharp, \text{unsafe}^\sharp$ : formula;
3  $P := \emptyset$ ;
4 while true do
5    $\text{unsafe}^\sharp := \text{lfp}(\text{pre}_P^\sharp, \text{bad})$ ;
6    $\text{safe}^\sharp := \text{lfp}(\text{pre}_P^\sharp, \text{final})$ ;
7   if  $(\text{safe}^\sharp \wedge \text{unsafe}^\sharp \wedge pc = \ell_0) \equiv \text{false}$  then
8     return  $(\text{safe}^\sharp \wedge pc = \ell_0)$ ;
9   Let  $\pi_E$  and  $\pi_F$  two traces s.t.  $\text{pre}_P^\sharp(\pi_E, \text{bad}) \wedge \text{pre}_P^\sharp(\pi_F, \text{final}) \not\equiv \text{false}$ ;
10  if  $\text{pre}(\pi_E, \text{bad}) \wedge \text{pre}(\pi_F, \text{final}) \not\equiv \text{false}$  then
11    print("warning: non-determinism in program");
12    exit;
13  else
14     $P := P \cup \text{Refine}(\pi_E, \pi_F)$ ;

```

---

**Proposition 1.** *The precondition  $\varphi$  computed by algorithm InferPrecond (a) guarantees the non-reachability of bad states and (b) the non-exclusion of safe terminating traces.*

*Proof.* **(a)**  $\varphi$  guarantees non-reachability of bad states. As computed by algorithm InferPrecond,  $\varphi \equiv \text{safe}^\sharp \wedge pc = \ell_0$ . Let us assume that there are states in  $\varphi$  from which a bad state can be reached. Thus, there is an error trace  $\pi_E$  such that

$$\text{pre}(\pi_E, \text{bad}) \wedge \text{safe}^\sharp \wedge pc = \ell_0 \not\equiv \text{false} \quad (5)$$

We also know that

$$\text{pre}(\pi_E, \text{bad}) \Rightarrow \text{unsafe}^\sharp, \quad (6)$$

as  $\text{lfp}(\text{pre}, \text{bad}) \leq \text{lfp}(\text{pre}^\sharp, \text{bad})$ . From (5) and (6) we obtain

$$\text{unsafe}^\sharp \wedge \text{safe}^\sharp \wedge pc = \ell_0 \not\equiv \text{false},$$

which contradicts the return condition at line 7 of algorithm InferPrecond.

**(b)**  $\varphi$  does not exclude safe terminating traces. Let us assume that  $\varphi$  excludes a given safe terminating trace  $\pi_F$  from  $\ell_0$  to  $\ell_F$ , which means that

$$\text{pre}(\pi_F, \text{final}) \wedge \text{safe}^\sharp \equiv \text{false}$$

or

$$\text{pre}(\pi_F, \text{final}) \Rightarrow \neg \text{safe}^\sharp. \quad (7)$$

We also have

$$\text{pre}(\pi_F, \text{final}) \Rightarrow \text{safe}^\sharp, \quad (8)$$

as  $\text{lfp}(\text{pre}, \text{final}) \leq \text{lfp}(\text{pre}^\sharp, \text{final})$ . From (7) and (8) we conclude  $\text{pre}(\pi_F, \text{final}) \equiv \text{false}$ , which means that such a trace  $\pi_F$  is not feasible.  $\square$

As program model checking is not decidable, we have no guarantee for termination of algorithm `InferPrecond`. However, whenever it terminates, the previous proposition holds.

## 4.2 Refinement for Precondition Inference

The main goal of refinement is to generate the minimal possible set of predicates that rule out a maximum number of spurious traces. Hence, the generated predicates must be as general as possible. We present procedure `Refine`, which takes as parameters two traces, one trace  $\pi_E$  leading to the error location, and another one  $\pi_F$  leading to the final location. The returned result is a set of predicates  $P$  that enables the verifier to show the following:

$$\text{pre}_P^\sharp(\pi_E, \text{bad}) \wedge \text{pre}_P^\sharp(\pi_F, \text{final}) \equiv \text{false} .$$

The procedure `Refine` relies on several other procedures: `atoms`, `MinCorePrio` and `ExtractNewPreds`. The procedure `atoms` is simply defined as

$$\text{atoms}(\varphi_1 \wedge \dots \wedge \varphi_n) = \{\varphi_1, \dots, \varphi_n\} .$$

It takes a conjunction  $\varphi$  and returns the set of its conjuncts.

The procedure `MinCorePrio` takes three arguments. The first one is a conjunction  $\varphi$ , the second one is an arbitrary formula  $\varphi'$  and the third one is a list  $L$  of formulas. As precondition,  $\varphi$  and  $\varphi'$  must be inconsistent. The procedure computes a minimal core of the conjunction  $\varphi$  that is inconsistent with the second argument  $\varphi'$ . There is usually more than one core that can be returned. This choice can be controlled by means of  $L$ , the third argument. `MinCorePrio` gives priority to the set of formulas in  $L$  to appear in the resulting minimal core, as illustrated by Algorithm 3. The list  $L$  is sorted in ascending order according to the priority of its elements. The algorithm proceeds by eliminating irrelevant predicates (conjuncts) of lesser priority (front) first. A predicate is irrelevant if its removal does not have an impact on the inconsistency of the new conjunction with  $\varphi'$ . The lowest priority is given to basic predicates in  $\varphi$  by storing them in the front of the list  $L$  (line 6). The consistency test at line 12 of the algorithm is carried out by calling a theorem prover.

Finally, procedure `ExtractNewPreds` implements a heuristic for predicate inference. It takes a conjunction as argument and returns a list of predicates sorted in ascending order of their likely importance to the convergence of the main CEGAR loop. We will describe this procedure in more details later in this section.

Back to the main procedure `Refine`, we see that it first computes the weakest precondition (`pre`) for each of the two traces taken as parameters (lines 4 and 5) to obtain formulas  $\psi_E$  and  $\psi_F$ . It then applies `ExtractNewPreds` to augment  $\psi_E$  and  $\psi_F$  with new facts induced by the two formulas (lines 6, 7, 9, 10). Finally, the minimal unsatisfiable cores of  $\psi_E$  and of  $\psi_F$  are computed (lines 8 and 11)

**Algorithm 2.** Refine

---

**Input:** two traces  $\pi_E$  and  $\pi_F$   
**Output:** set of predicates  $P$

- 1 **Var**  $P$  (initially empty),  $S_E, S_F$ : set of formulas;
- 2 **Var**  $P_{new}$ : list of formulas;
- 3 **Var**  $\psi_E, \psi_F$ : formula;
- 4  $\psi_E := \text{pre}(\pi_E, \text{bad})$ ;
- 5  $\psi_F := \text{pre}(\pi_F, \text{final})$ ;
- 6  $P_{new} := \text{ExtractNewPreds}(\psi_E)$ ;
- 7  $\psi_E := \psi_E \wedge (\bigwedge_{p \in P_{new}} p)$ ;
- 8  $\psi_E := \text{MinCorePrio}(\psi_E, \psi_F, P_{new})$ ;
- 9  $P_{new} := \text{ExtractNewPreds}(\psi_F)$ ;
- 10  $\psi_F := \psi_F \wedge (\bigwedge_{p \in P_{new}} p)$ ;
- 11  $\psi_F := \text{MinCorePrio}(\psi_F, \psi_E, P_{new})$ ;
- 12  $P := P \cup \text{atoms}(\psi_E) \cup \text{atoms}(\psi_F)$  ;
- 13 **Let**  $\pi_E = \tau_1; \dots; \tau_i$ ;
- 14 **Let**  $\pi_F = \tau'_1; \dots; \tau'_j$ ;
- 15 **Let**  $S_E = \bigcup_{k=1}^i \{\varphi_k\}$  s.t.  $\varphi_k \equiv \text{pre}(\tau_k; \dots; \tau_i, \text{bad})$ ;
- 16 **Let**  $S_F = \bigcup_{k=1}^j \{\varphi'_k\}$  s.t.  $\varphi'_k \equiv \text{pre}(\tau'_k; \dots; \tau'_j, \text{final})$ ;
- 17 **foreach**  $k$  in range  $\{1, \dots, i-1\}$  **do**
- 18      $P_{new} := \text{ExtractNewPreds}(\varphi_{k+1})$ ;
- 19      $\varphi_{k+1} := \varphi_{k+1} \wedge (\bigwedge_{p \in P_{new}} p)$ ;
- 20      $\psi_E := \text{MinCorePrio}(\varphi_{k+1}[X'/X], \tau_k \wedge \neg\psi_E, P_{new}[X'/X])$ ;
- 21      $P := P \cup \text{atoms}(\psi_E[X/X'])$  ;
- 22 **foreach**  $k$  in range  $\{1, \dots, j-1\}$  **do**
- 23      $P_{new} := \text{ExtractNewPreds}(\varphi'_{k+1})$ ;
- 24      $\varphi'_{k+1} := \varphi'_{k+1} \wedge (\bigwedge_{p \in P_{new}} p)$ ;
- 25      $\psi_F := \text{MinCorePrio}(\varphi'_{k+1}[X'/X], \tau'_k \wedge \neg\psi_F, P_{new}[X'/X])$ ;
- 26      $P := P \cup \text{atoms}(\psi_F[X/X'])$  ;
- 27 **return**  $P$ ;

---

and conjuncts appearing in either of them are added to the set of predicates (line 12).

In the next phase of the algorithm, the two formulas  $\psi_E$  and  $\psi_F$  are used to guide the inference of new predicates from states ( $\varphi_k$ 's and  $\varphi'_k$ 's) belonging to the error trace  $\pi_E$  (first loop, lines 17–21) and to the safe one  $\pi_F$  (second loop, lines 22–26). Along each trace and for each triple of pre-state  $\psi$ , transition  $\tau$  and post-state  $\varphi$ , we want to compute the minimal core  $\varphi_m$  of  $\varphi$  augmented with facts inferred via `ExtractNewPreds` such that  $\text{pre}(\tau, \varphi_m) \Rightarrow \psi$ , i.e.,  $\varphi_m[X'/X] \wedge \tau \wedge \neg\psi \equiv \text{false}$ . This amounts to computing the minimal core of  $\varphi_m[X'/X]$  with respect to  $\tau \wedge \neg\psi$ , as performed in lines 20 and 25 of the algorithm.

The procedure `ExtractNewPreds` is applied to the states of  $\pi_E$  and  $\pi_F$ , i.e., the  $\varphi_k$ 's and  $\varphi'_k$ 's of each trace. These states are obtained via a backward analysis of  $\pi_E$  and  $\pi_F$  during the initial phase of the algorithm (lines 4 and 5). As mentioned earlier, the operator `pre` (also `wp` in our case) is limited in inferring

**Algorithm 3.** MinCorePrio

---

**Input:**  $\varphi$  a conjunction of formulas,  $\varphi'$  a formula,  $L$  a list of formulas  
**Output:** a conjunction of formulas

- 1 **Var**  $\varphi''$ : formula;
- 2 **Var**  $S$ : set of formulas;
- 3 **Var**  $L, L'$ : list of formulas;
- 4  $S := \text{atoms}(\varphi)$ ;
- 5  $L' := L$ ;
- 6 add elements of  $S$  to  $L'$  in the front;
- 7 add elements of  $L$  to  $S$ ;
- 8 **foreach** formula  $\varphi_L \in L'$  **do**
- 9     **if**  $S - \{\varphi_L\} = \emptyset$  **then**
- 10     |     **return**  $\varphi_L$ ;
- 11      $\varphi'' := \bigwedge \varphi \mid \varphi \in S - \{\varphi_L\}$ ;
- 12     **if**  $\varphi'' \wedge \varphi' \equiv \text{false}$  **then**
- 13     |      $S := S - \{\varphi_L\}$ ;
- 14  $\varphi'' := \bigwedge \varphi \mid \varphi \in S$ ;
- 15 **return**  $\varphi''$ ;

---

relevant predicates, as it fails to generalize. Therefore, procedure `MinCorePrio` biases the computation of the minimal core by giving priority to predicates found via `ExtractNewPreds`, which are more likely to be general.

### 4.3 Predicate Inference

The procedure `ExtractNewPreds` plays a key role in our approach. It is based on a system of inference rules in the spirit of [23], where an interpolation procedure [19] is used to find predicates, followed by the application of a system of inference rules to deduce *range predicates*. In [23], the interpolant provides a concise description of the cause of infeasibility of traces, thus the base formula is already minimal. However, the application of the inference rules may introduce redundancies. In our case, `MinCorePrio` is applied after inferring the new facts, hence, it prevents the inundation of the system with irrelevant predicates. This is not just an optimization: during our experiments, this step has often made the difference between termination and divergence. The system of inference rules that we are using is given in Figure 4.

*Predicate inference system.* Divergence of the refinement process is often caused by predicates over variables that are increasing or decreasing (counters). This leads to the generation of sequences of constants when loops are effectively unfolded. Another cause of divergence are arrays with counter variables in their index expressions. A simple solution, advocated by [22], is to (initially) discard such predicates.

$$\begin{array}{c}
\frac{c_1.e + e_1 \geq 0, -c_2.e + e_2 \geq 0}{\frac{c_2.e_1 + c_1.e_2 \geq 0}{(c_1, c_2 > 0)}} \text{ (ELIM)} \qquad \frac{x - e \geq 0, -x + e \geq 0}{x = e} \text{ (EQ)} \\
\\
\frac{\varphi(x), x = e}{\varphi(e)} \text{ (SUB)} \qquad \frac{\varphi(i), \neg\varphi(j) \ (i < j)}{\exists x \in \{i, \dots, j\}. \varphi(x), \exists x \in \{i, \dots, j\}. \neg\varphi(x)} \text{ (EXIST)} \\
\\
\frac{\exists x \in \{i, \dots, j\}. \varphi(x), j \leq k}{\exists x \in \{i, \dots, k\}. \varphi(x)} \text{ (EXT\_R)} \qquad \frac{\exists x \in \{i, \dots, j\}. \varphi(x), k \leq i}{\exists x \in \{k, \dots, j\}. \varphi(x)} \text{ (EXT\_L)} \\
\\
\frac{\varphi(i)}{\forall x \in \{i\}. \varphi(x)} \text{ (UNIV)} \qquad \frac{\forall x \in \{j, \dots, i\}. \varphi(x), \forall x \in \{i+1, \dots, k\}. \varphi(x)}{\forall x \in \{j, \dots, k\}. \varphi(x)} \text{ (LINK)}
\end{array}$$

*i* and *j* are integer variables appearing in a linear index expression in  $\varphi$  ( $\neg\varphi$ ).

**Fig. 4.** Rules for predicate inference

The aim of the system of rules of Figure 4 is to eliminate likely diverging sequences of predicates whenever possible by inferring new predicates that are more general. Among the symbols used in the system,  $e$  refers to linear terms,  $x$  is a variable and  $\varphi$  is a formula. The rule ELIM linearly combines two constraints to eliminate common variables. Rule EQ infers equality constraints, which might be used by rule SUB to substitute occurrences of variables with equal terms. The rule UNIV builds a quantified formula and LINK bridges the intervals of two quantified formulas. Finally, the rule EXIST produces two existentially quantified formulas and the rules EXT\_R and EXT\_L extend the interval of an existentially quantified formula from the right and the left, respectively.

The procedure `ExtractNewPreds` (Algorithm 4) applies the rules of the inference system to the conjuncts of the formula given as argument and returns a list of predicates sorted in ascending order of priority. A predicate  $p_1$  has higher priority than predicate  $p_2$  if  $p_1$  is produced by a rule where  $p_2$  appears as one of its antecedents. The procedure starts with the list of basic predicates that are extracted from the formula given as argument. These predicates have the lowest priority. It then keeps applying the rules to predicates in the list until saturation, i.e., until no new predicates are produced. The code fragment from line 13 to 17 stores the new predicates according to their priority, i.e., in a position of the list that is beyond the positions of the associated antecedents.

The algorithm terminates as the two rules ELIM and EQ are only applied to basic predicates (condition at line 12). Thus, they will be called a finite number of times generating a finite number of linear constraints. All other rules will generate a finite number of predicates, as they all depend on linear constraints. We furthermore do not consider nested array expressions. The order in which the rules are applied does not matter.

*Illustration.* Let us illustrate the application of procedure `Refine` to program `copy_2` of Figure 1(b). We call `Refine` with the error trace  $\langle \ell_0, \ell_1, \ell_2, \ell_1, \ell_2, \ell_E \rangle$  and the safe trace  $\langle \ell_0, \ell_1, \ell_2, \ell_1, \ell_F \rangle$ , which both enter the loop in program `copy_2`

**Algorithm 4.** ExtractNewPreds

---

```

Input: formula  $\varphi$ 
Output: list of formulas
1 Var  $S, S_b$ : set of formulas;
2 Var  $L, L'$ : list of formulas;
3 Var  $R$ : list of inference rules;
4  $S_b := \text{atoms}(\varphi)$ ;
5 insert elements of  $S_b$  in  $L'$ ;
6  $R := \{\text{ELIM}, \text{EQ}, \text{UNIV}, \text{SUB}, \text{LINK}, \text{EXIST}, \text{EXT\_L}, \text{EXT\_R}\}$ ;
7 repeat
8    $L := L'$ ;
9   foreach rule  $r \in R$  do
10     Let  $k$  be the number of premises of  $r$ ;
11     foreach tuple  $t \in L^k$  do
12       if ( $r \notin \{\text{ELIM}, \text{EQ}\}$ )  $\vee (\forall i \in \{1, \dots, k\}. t_i \in S_b)$  then
13          $S := r(t)$ ;
14         Let  $pos = \max\{pos_j \mid j \in \{1, \dots, k\} \wedge L[pos_j] = t_j\}$ ;
15         foreach predicate  $p \in S$  do
16           if  $p \notin L'$  then
17              $\lfloor$  insert  $p$  after position  $pos$  in  $L'$ ;
18 until  $L = L'$ ;
19 return  $L$ ;
```

---

once. The analysis of these two traces is illustrated in Figure 5. The upper table shows results for the error trace and the lower one for the safe trace. In both tables, the first column contains the suffix of the trace that is analyzed backwards using the weakest precondition. The result is shown in the second column. Finally, the third column shows the new predicates that are inferred using the information from the second column. The superscript associated with each predicate is its priority. At the initial location, which corresponds to the second line in both tables, the predicate  $\forall x \in \{0, \dots, a\}. b[x] \neq 0$  is the one with the highest priority for the error trace. It is inferred via the application of the rule UNIV followed by SUB.

For the safe trace, we have two predicates of the highest priority, namely  $\exists x \in \{0, \dots, a\}. b[x] \neq 0$  and  $\exists x \in \{0, \dots, a\}. b[x] = 0$ . They are both generated by applying rules EXIST and EXT\_R successively. The refinement procedure selects the second predicate as it is the one which separates the two initial states. The selected predicates are underlined in both tables. Going one step backward from the initial location  $\ell_0$  to location  $\ell_1$  in both traces, the selected predicates are  $\forall x \in \{i, \dots, a\}. b[x] \neq 0$  and  $\exists x \in \{i, \dots, a\}. b[x] = 0$  for the error and safe trace, respectively. These are the predicates on which the ones selected at the initial location  $\ell_0$  depend. One can assert that these two predicates are backwards invariants with respect to the cycle  $\langle \ell_1, \ell_2, \ell_1 \rangle$ . They thus cover an infinite number of traces.

Error trace	WP	New predicates
$\ell_1, \ell_2, \ell_1, \ell_2, \ell_E$	$i + 1 \geq a\llcorner, b[i + 1] \neq 0, i < a\lrcorner$ $b[i] \neq 0$	$a\llcorner = i + 1^{(1)}, \forall x \in [i, i + 1]. b[x] \neq 0^{(1)}$ <u><math>\forall x \in [i, a\llcorner]. b[x] \neq 0^{(2)}</math></u>
$\ell_0, \ell_1, \ell_2, \ell_1, \ell_2, \ell_E$	$1 \geq a\llcorner, b[1] \neq 0, 0 < a\lrcorner$ $b[0] \neq 0$	$a\llcorner = 1^{(1)}, \forall x \in [0, 1]. b[x] \neq 0^{(1)}$ <u><math>\forall x \in [0, a\llcorner]. b[x] \neq 0^{(2)}</math></u>

Safe trace	WP	New predicates
$\ell_1, \ell_2, \ell_1, \ell_F$	$b[i] \neq 0, i + 1 \leq a\llcorner, b[i + 1] = 0$	$\exists x \in [i, i + 1]. b[x] \neq 0^{(1)}, \exists x \in [i, i + 1]. b[x] = 0^{(1)}$ <u><math>\exists x \in [i, a\llcorner]. b[x] \neq 0^{(2)}, \exists x \in [i, a\llcorner]. b[x] = 0^{(2)}</math></u>
$\ell_0, \ell_1, \ell_2, \ell_1, \ell_F$	$b[0] \neq 0, 1 \leq a\llcorner, b[1] = 0$	$\exists x \in [0, 1]. b[x] \neq 0^{(1)}, \exists x \in [0, 1]. b[x] = 0^{(1)}$ <u><math>\exists x \in [0, a\llcorner]. b[x] \neq 0^{(2)}, \exists x \in [0, a\llcorner]. b[x] = 0^{(2)}</math></u>

**Fig. 5.** Illustration of algorithm `Refine` on program `copy_2`. The underlined predicates are selected by the refinement process. The superscript is the priority of each predicate.

## 5 Experimental Results

*Implementation.* We have implemented our precondition inference technique in the P-Gen<sup>4</sup> tool. P-Gen takes as input a C program containing a procedure annotated with an assertion to be verified. As output, it returns a formula that represents the set of pre-states from which the specified assertion holds for any execution.

*Experiments.* We performed experiments using a desktop computer with 3.7 GB of RAM and a quad-core processor with 2.83 GHz, running Linux. P-Gen uses several theorem provers to compute the abstraction and analyze counterexamples. We have initially used Yices [15] and Simplify [14], but observed limitations when handling quantified formulas. These limitations often lead to the divergence of CEGAR, as the refinement procedure picks up a set of quantifier-free predicates instead of a quantified predicate, and thus fails to generalize. We have subsequently integrated Z3 [13] and used it running as a standalone process communicating with P-Gen through pipes. The Z3 theorem prover was able to decide many queries that were not handled by the two other theorem provers.

The results of our experiments are summarized in Table 1. The column “Precond.” shows the type of precondition inferred (“Q” stands for quantified and “S” stands for simple, i.e., quantifier-free). The column “Iter.U.” (“Iter.S.”) gives the number of iterations performed by CEGAR to compute the set of unsafe (safe) states and column “Pred.U.” (“Pred.S.”) gives the number of predicates inferred to abstract the set of unsafe (safe) states. Our tool is based on lazy abstraction [20], we therefore provide the average number of predicates per location instead of the total number of predicates. This number is an indicator for memory consumption, as predicates are encoding program states. Finally, column “Alt.” refers to the number of alternations between the two abstractions

<sup>4</sup> <http://www.cs.ox.ac.uk/people/nassim.seghir/pgen-web-page>

**Table 1.** Experimental results for routines taken from the C string library (upper part) and from real-world programs (lower part)

Program	Precond.	Iter.U.	Pred.U.	Iter.S.	Pred.S.	Alt.	Time (s)
memcmp	Q + S	5	5	8	3	2	23.64
strcat	Q	4	4	4	3	2	0.77
memchr	Q + S	5	4	8	3	2	77.30
strlen	Q	4	4	4	3	2	0.80
memcpy	S	3	2	3	2	1	0.17
memmove	S	5	4	9	2	2	0.91
strchr	Q	5	7	7	4	2	1.92
r_strcat	Q	5	2	3	2	2	9.02
strcspn	Q	5	6	7	4	2	7.30
strspn	Q	5	6	7	4	2	7.05
my_strcmp	Q	5	7	7	4	2	2.70
my_strcpy	Q	4	4	4	3	2	1.41
AllNotNull	Q + S	6	5	5	3	2	2.07
perfc_copy_info	S	7	2	13	2	2	1.94
bitmap_shift_right	S	15	13	26	5	2	41.60
mvswap	S	7	5	8	4	3	0.70
BZ2_hbAssignCodes	S	8	4	6	4	2	0.73

(i.e., the number of times the procedure switches between the abstraction of unsafe states and the abstraction of safe states).

The upper part of the table relates to implementations of routines from the C string library<sup>5</sup>. The last two procedures, `my_strcmp` and `my_strcpy`, are modified versions of the original `strcmp` and `strcpy`. In the lower part of the table, we have `AllNotNull`, which was used by Cousot et al. as illustration [12]. Our tool infers a precondition equivalent to the one proposed in their paper. The procedures `mvswap` and `BZ2_hbAssignCodes` are taken from the source code of the compression program `Bzip2`<sup>6</sup>. The procedures `perfc_copy_info` and `bitmap_shift_right` belong to the Xen Hypervisor<sup>7</sup>.

Although the refinement heuristic performs well in most of the cases, sometimes it does not select adequate predicates in early stages of the iterative process (as for `memchr`, `memcmp` and `bitmap_shift_right`), resulting in high time consumption. However, due to the precision of the precondition, this cost can often be amortised when checking code that uses the functions.

## 6 Related Work

The work presented in this paper is linked to several topics: predicate abstraction, invariant generation, counterexample-guided refinement, modular verification and

<sup>5</sup> [http://en.wikibooks.org/wiki/C\\_Programming/Strings](http://en.wikibooks.org/wiki/C_Programming/Strings).

<sup>6</sup> <http://www.bzip.org/>

<sup>7</sup> <http://www.xen.org/products/xenhyp.html>

precondition inference. Along these axes, we elaborate on some related work from the literature.

The combination of predicate abstraction [16] with counterexample-guided abstraction refinement [9] has been pioneered by the SLAM [1] tool at Microsoft, and has subsequently been implemented in many other tools, including BLAST [20], MAGIC [7], ARMC [25], F-SOFT [21] and SATABS [8]. The DASH [2] and SYNERGY [17] algorithms are variants of this approach, as they are based on program testing to refine abstractions. Even closer to our approach, the tool YASM [18] combines over- and under-approximations to prove or refute program properties. However, all these tools and methods aim at checking the validity of a given assertion. In contrast, we use CEGAR to compute a precondition under which the assertion is valid. Our technique represents a general scheme for extending the previously mentioned tools to infer preconditions, as it uses ingredients that are common to all of them.

Moy proposed a technique to infer preconditions [24] using a combination of state-of-the-art techniques such as abstract interpretation, weakest preconditions and quantifier elimination. While his technique is stronger than many existing ones, it is unable to infer quantified preconditions. Our technique is able to infer universally- as well as existentially-quantified preconditions. Blanc and Kroening proposed a technique to optimize the simulation of SystemC code [3]. It consists of inferring conditions that are sufficient for the commutativity of pairs of processes. Like our algorithm, theirs also is based on CEGAR. However, they have no guarantee that the inferred precondition is exact. Our method does provide this guarantee.

Taghdiri proposed an approach for generating approximations of relations (over pre- and post-states) induced by functions [28]. A function specification (relation) is computed with respect to a context that includes the property to be verified, thus the specification may lack generality. Moreover, the number of times loops are unrolled is bounded, making the approach unsuitable for proving the absence of bugs. Our technique over-approximates the set of all behaviors. Thus, the preconditions computed by our method guarantee safety.

Sankaranarayanan et al. presented a technique that combines test and machine learning to infer likely data preconditions over a set of predicates [27]. In many cases, they were able to learn preconditions that ensure safe executions. However, as their technique is based on testing, it can only suggest preconditions, but does not guarantee their soundness.

In the context of abstract interpretation, Cousot et al. formulated the contract inference problem for intermittent assertions precisely [12]. The precondition extracted by their method does not exclude safe runs even when a non-deterministic choice could lead to bad ones. Our treatment of non-determinism is different, as we report a warning to the user. The method described in [12] as well as [4] and [26] rely on some predefined abstract domains. In our approach, the precision of the abstraction is automatically tuned as required by means of CEGAR. We can also enhance the refinement process by introducing new inference rules, without having to implement new transformers.

Calcagno et al. presented a technique based on Bi-Abduction to infer pre- and post-specifications of the heap [6]. One of the major advantages of their approach is scalability. Similar to their approach, we intend to complete the picture by integrating our method into an inter-procedural reasoning framework. Although we can deal with pointers, the properties handled by their technique are out of the scope for our tool as we do not have a theory to reason about heap properties. On the other side, contrary to our approach, the preconditions they compute are not exact, meaning that they might have to be refined. In the context of termination, Bozga et al. [5], proposed a method to generate exact preconditions for a restricted class of programs. Other techniques for inferring preconditions for termination are applicable to a larger set of programs [10], however they only generate sufficient preconditions.

## 7 Conclusion

We have presented a new method for precondition inference based on counterexample-guided abstraction refinement. Given a procedure containing an assertion, our method infers a formula that is sufficient for the validity of the specified assertion and exclusively refers to state variables visible at the entry point of the procedure. The inferred precondition is independent of the context, making it reusable. The computed precondition is also necessary for the validity of the assertion, i.e., it does not exclude any safe runs of the procedure. Hence, we avoid false alarms, as the violation of the precondition corresponds to a real error.

Our technique is based on ingredients commonly used in CEGAR-based assertion checkers, and it thus represents a general scheme for extending other verification tools to infer preconditions. Moreover, we believe that we can take advantage of components from other tools, such as advanced refinement mechanisms, which can be integrated seamlessly into our algorithm. Preliminary experimental results are encouraging, as we are able to generate exact preconditions (quantified as well as quantifier-free) for string- and array-manipulating programs.

## References

1. Ball, T., Rajamani, S.K.: The SLAM project: debugging system software via static analysis. In: POPL, pp. 1–3 (2002)
2. Beckman, N.E., Nori, A.V., Rajamani, S.K., Simmons, R.J.: Proofs from tests. In: ISSA, pp. 3–14 (2008)
3. Blanc, N., Kroening, D.: Race analysis for SystemC using model checking. In: Proceedings of ICCAD 2008, pp. 356–363. IEEE (2008)
4. Bourdoncle, F.: Abstract debugging of higher-order imperative languages. In: PLDI, pp. 46–55 (1993)
5. Bozga, M., Iosif, R., Konečný, F.: Deciding Conditional Termination. In: Flanagan, C., König, B. (eds.) TACAS 2012. LNCS, vol. 7214, pp. 252–266. Springer, Heidelberg (2012)

6. Calcagno, C., Distefano, D., O'Hearn, P.W., Yang, H.: Compositional shape analysis by means of bi-abduction. In: POPL, pp. 289–300 (2009)
7. Chaki, S., Clarke, E.M., Groce, A., Jha, S., Veith, H.: Modular verification of software components in C. In: ICSE, pp. 385–395 (2003)
8. Clarke, E., Kroening, D., Sharygina, N., Yorav, K.: Predicate abstraction of ANSI-C programs using SAT. *Formal Methods in System Design (FMSD)* 25, 105–127 (2004)
9. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-Guided Abstraction Refinement. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 154–169. Springer, Heidelberg (2000)
10. Cook, B., Gulwani, S., Lev-Ami, T., Rybalchenko, A., Sagiv, M.: Proving Conditional Termination. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 328–340. Springer, Heidelberg (2008)
11. Cook, B., Podelski, A., Rybalchenko, A.: TERMINATOR: Beyond Safety. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 415–418. Springer, Heidelberg (2006)
12. Cousot, P., Cousot, R., Logozzo, F.: Precondition Inference from Intermittent Assertions and Application to Contracts on Collections. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 150–168. Springer, Heidelberg (2011)
13. de Moura, L., Bjørner, N.: Z3: An Efficient SMT Solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
14. Detlefs, D., Nelson, G., Saxe, J.B.: Simplify: A theorem prover for program checking. Technical Report HPL-2003-148, HP Lab. (2003)
15. Dutertre, B., de Moura, L.: A Fast Linear-Arithmetic Solver for DPLL(T). In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 81–94. Springer, Heidelberg (2006)
16. Graf, S., Saïdi, H.: Construction of Abstract State Graphs with PVS. In: Grumberg, O. (ed.) CAV 1997. LNCS, vol. 1254, pp. 72–83. Springer, Heidelberg (1997)
17. Gulavani, B.S., Henzinger, T.A., Kannan, Y., Nori, A.V., Rajamani, S.K.: SYN-ERGY: a new algorithm for property checking. In: SIGSOFT FSE, pp. 117–127 (2006)
18. Gurfinkel, A., Wei, O., Chechik, M.: YASM: A Software Model-Checker for Verification and Refutation. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 170–174. Springer, Heidelberg (2006)
19. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. In: POPL, pp. 232–244 (2004)
20. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: POPL, pp. 58–70 (2002)
21. Ivancic, F., Shlyakhter, I., Gupta, A., Ganai, M.K.: Model checking C programs using F-SOFT. In: ICCD, pp. 297–308 (2005)
22. Jhala, R., McMillan, K.L.: A Practical and Complete Approach to Predicate Refinement. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 459–473. Springer, Heidelberg (2006)
23. Jhala, R., McMillan, K.L.: Array Abstractions from Proofs. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 193–206. Springer, Heidelberg (2007)

24. Moy, Y.: Sufficient Preconditions for Modular Assertion Checking. In: Logozzo, F., Peled, D.A., Zuck, L.D. (eds.) VMCAI 2008. LNCS, vol. 4905, pp. 188–202. Springer, Heidelberg (2008)
25. Podelski, A., Rybalchenko, A.: ARMC: The Logical Choice for Software Model Checking with Abstraction Refinement. In: Hanus, M. (ed.) PADL 2007. LNCS, vol. 4354, pp. 245–259. Springer, Heidelberg (2007)
26. Rival, X.: Understanding the Origin of Alarms in ASTRÉE. In: Hankin, C., Siveroni, I. (eds.) SAS 2005. LNCS, vol. 3672, pp. 303–319. Springer, Heidelberg (2005)
27. Sankaranarayanan, S., Chaudhuri, S., Ivancic, F., Gupta, A.: Dynamic inference of likely data preconditions over predicates by tree learning. In: ISSTA, pp. 295–306 (2008)
28. Taghdiri, M.: Inferring specifications to detect errors in code. In: ASE, pp. 144–153 (2004)