

Explaining the Effectiveness of Small Refinement Heuristics in Program Verification with CEGAR

Tachio Terauchi

JAIST

terauchi@jaist.ac.jp

Abstract. Safety property (i.e., reachability) verification is undecidable for Turing-complete programming languages. Nonetheless, recent progress has led to heuristic verification methods, such as the ones based on predicate abstraction with counterexample-guided refinement (CEGAR), that work surprisingly well in practice. In this paper, we investigate the effectiveness of the *small refinement heuristic* which, for abstraction refinement in CEGAR, uses (the predicates in) a small proof of the given counterexample’s spuriousness [12, 3, 17, 22]. Such a method has shown to be quite effective in practice but thus far lacked a theoretical backing. Our core result is that the heuristic guarantees certain bounds on the number of CEGAR iterations, relative to the size of a proof for the input program.

1 Introduction

The safety property (i.e., reachability) verification problem asks, given a program, if an error state (e.g., given as a line number in the program text) is unreachable for every execution of the program. The problem is undecidable for Turing-complete programming languages (and intractably hard for many natural decidable fragments – e.g., PSPACE-complete for Boolean programs). Despite the staggering complexity, recent research has led to heuristic verification methods that work surprisingly well in practice. They have been used to verify non-trivial real-world programs such as operating system device drivers [4, 11], and the yearly held software verification competition [1] shows an ever increasing variety of programs efficiently verified by the state-of-the-art automated software verifiers. By contrast, there has been comparatively less progress on explaining why such heuristics work. As a step toward bridging the gap, we present a theoretical explanation for the effectiveness of a heuristic used in predicate abstraction with counterexample-guided refinement (CEGAR).

CEGAR is a verification method that iteratively updates a finite set of predicates from some first-order logic (FOL) theory, called the *candidate predicate set*, until the candidate set forms a sufficient proof of the given program’s safety (i.e., an inductive invariant). In each iteration, a process called *abstraction* checks if the current candidate set is sufficient, and if not, a counterexample is generated. A process called *refinement* infers a proof of the counterexample’s safety (i.e.,

spuriousness) whereby the predicates in the proof are added as new candidates, and the iteration repeats.

In general, there is more than one proof that refutes the same counterexample, and which proof is inferred by the refinement process can significantly affect the performance of the overall verification process (cf. Section 2.2 for an example). A heuristic often used in practice is to infer a *small* proof (e.g., measured by the sum of the syntactic size of the predicates), and researchers have proposed methods for inferring a small proof of the given counterexample’s spuriousness [12, 3, 17, 22].¹ This paper analyzes the effect such a *small refinement heuristic* has on the overall verification performance.

As safety property verification is undecidable, we cannot establish any (finitary) complexity result without some assumption on the input problem instances. In general, we will state our results relative to the size of a proof of safety for the input program, and we will try to show that the verification converges quickly assuming that the input program has a small proof (e.g., polynomial in the size of the program). The assumption captures the conventional wisdom that correct programs are often correct for simple reasons, per Occam’s razor.

Overview of the Main Results. We formalize the small refinement heuristic to be a refinement process that infers a proof of size at most polynomial in that of the smallest proof for the given counterexample. Let $\text{CEGVERIF}_{\text{SR}}$ be a CEGAR verification with such a refinement process. Let $\text{mpfsize}(P)$ be the size of the smallest proof of safety for a program P . Our first main result is the following.

Theorem 1. $\text{CEGVERIF}_{\text{SR}}$ converges in at most $\exp(\text{mpfsize}(P))$ many CEGAR iterations given a program P .

Theorem 1 implies that $\text{CEGVERIF}_{\text{SR}}$ is able to verify a program in an exponential number of iterations under the promise that the program has a polynomial size proof of safety. We prove Theorem 1 under a rather general setting.

Next, we consider a more concrete setting where a program is represented by a control flow graph (CFG) such that the program’s proof is a Floyd-style node-wise inductive invariant [9], the abstraction process uses the Cartesian predicate abstraction, and the size of a proof is measured by the sum of the syntactic size of the predicates. We also assume that counterexamples are generated by unfolding each loop an arbitrary but the same number of times (copies of nested subloops are also unfolded the same number of times) which we show to be sufficient for the setting (cf. Theorem 9). Under such a setting, we show that $\text{CEGVERIF}_{\text{SR}}$ converges in at most $\text{poly}(\text{mpfsize}(P))^{\text{maxhds}(P)}$ many iterations where $\text{maxhds}(P)$ is the maximum number of loop entries per loop of the CFG (e.g., $\text{maxhds}(P) \leq 1$ for reducible CFGs) (Theorem 10). The result implies that, under such a setting, $\text{CEGVERIF}_{\text{SR}}$ is able to verify a program with a constant number of loop entries per loop in a polynomial number of iterations under the promise that the program has a polynomial size proof of safety.

¹ In this paper, a “proof” is a set of predicates. Note that this differs from the notion of proofs and proof sizes used in proof complexity [6].

The rest of the paper is organized as follows. We discuss related work next. Section 2 gives preliminary definitions pertaining to generic CEGAR verification and proves the first main result (Theorem 1). Section 3 gives additional preliminary definitions pertaining to CFG programs and proves the main result on CFG programs (Theorem 10). We discuss some limitations of our results in Section 4, and conclude the paper in Section 5. Appendix contains the proofs and extra materials omitted from the main body of the paper.

Related Work. Previous works on the small refinement heuristic [12, 3, 17, 22] have presented empirical evidences of the heuristic’s effectiveness in the form of experiments. To our knowledge, this paper is the first work on a theoretical explanation for the heuristic’s effectiveness.

A related but somewhat different heuristic proposed for CEGAR is the *stratified refinement* approach [13, 15, 20] where the proofs inferred in each refinement step is restricted to some finite set of proofs that is enlarged when the refinement fails to find a proof in the current set. The proof set is enlarged in such a way to ensure that the growing strata of proof sets eventually cover the underlying (possibly infinite) space of possible proofs. The stratified approach ensures an eventual convergence of CEGAR iterations under the promise that a proof of the program’s safety exists in the underlying proof space, but the previous works give no results on the iteration bound.

2 Iteration Bound Under a Generic Setting

2.1 Preliminaries

Formulas and Predicates. We write finite sequences in boldface (e.g., \mathbf{x} for x_1, \dots, x_n). Let \mathcal{T} be a FOL theory. A *term* a and *formula* ϕ in the signature of \mathcal{T} is defined as follows.

$$\begin{aligned} \text{term } a &::= x \mid f(\mathbf{a}) \\ \text{formula } \phi &::= p(\mathbf{a}) \mid \neg\phi \mid \phi \wedge \phi' \mid \phi \vee \phi' \mid \phi \Rightarrow \phi' \mid \exists x.\phi \mid \forall x.\phi \end{aligned}$$

Here, x is a variable, f is an arity $|\mathbf{a}|$ function symbol, and p is an arity $|\mathbf{a}|$ predicate symbol of \mathcal{T} . As usual, \Rightarrow binds weaker than \wedge or \vee . For a formula ϕ , we write $fv(\phi)$ for the set of free variables in ϕ . A *predicate* in \mathcal{T} is of the form $\lambda\mathbf{x}.\phi$ where ϕ is a formula such that $\{\mathbf{x}\} = fv(\phi)$. We often omit the explicit λ abstraction and treat a formula ϕ as the predicate $\lambda\mathbf{x}.\phi$ where $\{\mathbf{x}\} = fv(\phi)$. We overload \mathcal{T} for the set of predicates in \mathcal{T} , and write $\mathcal{T}(\mathbf{x})$ for the set of \mathcal{T} -predicates of arity $|\mathbf{x}|$. We write $\models_{\mathcal{T}} \phi$ if ϕ is valid in \mathcal{T} . We write \top for tautology and \perp for contradiction. (Note that the term “predicate” is not limited to just atomic predicates as sometimes is in the literature on CEGAR.)

Generic CEGAR. To state Theorem 1 for a general setting, following the style of [20], we give a definition of CEGAR in terms of generic properties that abstraction and refinement processes are assumed to satisfy, but without specifying exactly when predicates form a proof of a program or how counterexamples are

```

1: procedure CEGVERIF( $P$ )
2:    $Cands := \emptyset$ ;
3:   repeat
4:     match Abs( $P, Cands$ ) with
5:       safe  $\rightarrow$  return safe
6:       |  $\pi \rightarrow$  match Ref( $\pi$ ) with
7:         unsafe  $\rightarrow$  return unsafe
8:         |  $F \rightarrow Cands := Cands \cup F$ 
9:   end repeat

```

Fig. 1. CEGVERIF

generated and refuted (we concretize such notions for CFG programs in Section 3.1).

In general, a *counterexample* is an “unwound” slice of the given program and is itself considered a program (typically, without loops or recursion). Let γ range over programs and counterexamples. For a finite set of predicates $F \subseteq \mathcal{T}$, we write $F \vdash \gamma$ to mean that F is a proof of safety of γ . We often simply say that F is a *proof of γ* when $F \vdash \gamma$ (i.e., omitting “of safety”), and if γ is a counterexample in addition, we sometimes say that F *refutes γ* , or F is a proof of *spuriousness* of γ . We require the proof relation \vdash to be monotonic on the predicate set, that is, if $F \vdash \gamma$ and $F \subseteq F'$, then $F' \vdash \gamma$ (i.e., having more predicates can only increase the ability to prove). We use F, F', F_1 , etc. to range over finite sets of predicates.

Figure 1 shows the overview of the verification process. CEGVERIF takes as input the program P to be verified, and initializes the candidate predicate set $Cands$ to \emptyset (line 2). Then, it repeats the abstract-and-refine loop (lines 4-8) until convergence. Abs is the abstraction process which takes as input a program and a finite set of \mathcal{T} -predicates. For a program P and a finite $F \subseteq \mathcal{T}$, Abs(P, F) either returns **safe**, indicating that P has been proved safe using the predicates from F , or returns a counterexample. In the former case, the verification process halts, returning **safe**. Ref is the refinement process, which, given a counterexample π , either returns a proof of π 's safety, or detects that π is irrefutable and returns **unsafe**. In the former case, the proof is added to the candidates, and in latter case, the verification halts by returning **unsafe**. For a run of CEGVERIF, the *number of CEGAR iterations* is defined to be the number of times the abstract-and-refine loop (lines 4-8) iterated.

We state the required assumptions on the abstraction and refinement processes. We require that if $F \vdash \gamma$ then Abs(γ, F) = **safe**, and that if Abs(P, F) returns a counterexample π then Abs(π, F) \neq **safe** (i.e., Abs proves the safety of the program given a sufficient set of predicates, and otherwise returns a counterexample that it cannot prove safe with the given predicates). We also require that if $F \vdash P$ and Abs(P, F') returns π , then $F \vdash \pi$ (i.e., a proof for a program is also a proof for any counterexample of the program). We require that Abs is *sound* in that it only proves safe programs safe, that is, Abs(γ, F') = **safe** only if

$\exists F \subseteq \mathcal{T}.F \vdash \gamma$. For the refinement process `Ref`, we require that `Ref(π)` returns F only if $F \vdash \pi$ (i.e., the returned proof is actually a proof for the counterexample), and that `Ref(π)` returns `unsafe` only if $\forall F \subseteq \mathcal{T}.F \not\vdash \pi$ (i.e., only irrefutable counterexamples are detected irrefutable). Finally, we require `Abs` and `Ref` to halt on all inputs. We note that these assumptions are quite weak and satisfied by virtually any CEGAR verifiers.²

It is easy to see that `CEGVERIF` is sound in that it only proves safe programs safe.

Theorem 2 (Soundness). *`CEGVERIF(P)` returns safe only if $\exists F \subseteq \mathcal{T}.F \vdash P$.*

Note that Theorem 2 says nothing about how fast (or whether) `CEGVERIF` converges. The main results of the paper (Theorems 1 and 10) show that, when `Ref` is made to return a small proof of the given counterexample, `CEGVERIF` is guaranteed to converge in a number of iterations bounded by the size of a proof for the given program.

It is also easy to see that `CEGVERIF` is “complete” in the sense that it only detects unprovable programs unprovable.

Theorem 3 (Completeness). *`CEGVERIF(P)` returns unsafe only if $\forall F \subseteq \mathcal{T}.F \not\vdash P$.*

Since the paper is only concerned with analyzing the behavior of CEGAR when given a provably safe program, in what follows, we disregard the situation where the given program is unprovable with the predicates from the background theory.

2.2 Main Result

To demonstrate the usefulness of the small refinement heuristic, we start with an example on which `CEGVERIF` (without the heuristic) may fail to converge despite the program having a small proof of safety (taken from [20]). Figure 2 shows the program P_{ex} . Here, `ndet()` returns a non-deterministic integer. The goal is to verify that the assertion failure is unreachable, that is, $a = b \Rightarrow y = x$ whenever line 9 is reached. We define a proof of the program to be the set of predicates that can be used as loop invariants at each of the loop heads (lines 3 and 6) and are sufficient to prove the unreachability (i.e., “safe” inductive invariant). For example, a possible proof is the singleton set $F_{inv} = \{a = b \Rightarrow y = x + z\}$.

```

1: int a = ndet(); int b = ndet();
2: int x = a; int y = b; int z = 0;
3: while (ndet()) {
4:   y++;z++;
5: }
6: while (z != 0) {
7:   y--;z--;
8: }
9: assert (a!=b || y=x);

```

Fig. 2. P_{ex}

² We do not impose $F \vdash P \Leftrightarrow \text{Abs}(P, F) = \text{safe}$ to allow modeling verifiers whose abstraction process can prove more from the same predicates than the refinement process (cf. Section 3.1).

<pre> 1: int a=ndet(); int b=ndet(); 2: int x=a; int y=b; int z=0; 3: if (ndet()) { 4: y++;z++; 5: } 6: if (z!=0) { 7: y--;z--; 8: } 9: assert (a!=b y=x); </pre>	<pre> 1: int a=ndet(); int b=ndet(); 2: int x=a; int y=b; int z=0; 3: if (ndet()) { 4: y++;z++; 5: } 6: if (ndet()) { 7: y++;z++; 8: } 9: if (z!=0) { 10: y--;z--; 11: } 12: if (z!=0) { 13: y--;z--; 14: } 15: assert (a!=b y=x); </pre>
--	--

π_1

π_2

Fig. 3. Counterexamples of P_{ex}

Running CEGVERIF on the program, **Abs** may return the counterexample π_1 shown in Figure 3 in the first iteration, obtained by unfolding each loop once. Viewing the unfolded **if** statements at lines 3 and 6 as one-iteration loops, it can be seen that F_{inv} is a proof of π_1 's safety. Thus, **Ref**(π_1) may return F_{inv} , which is also a proof for P_{ex} and would allow CEGVERIF to converge in the next iteration.

Unfortunately, the refinement process is not guaranteed to return F_{inv} but may choose any set of predicates that forms a proof of π_1 's safety. For example, another possibility is $F_1 = \{\phi_0, \phi_1, \phi_0 \vee \phi_1\}$ where

$$\begin{aligned} \phi_0 &\equiv x = a \wedge y = b \wedge z = 0 \\ \phi_1 &\equiv x = a \wedge y = b + 1 \wedge z = 1 \end{aligned}$$

Adding F_1 to the candidates is sufficient for proving the safety of π_1 but not that of P_{ex} , and so the abstraction process in the subsequent iteration would return yet another counterexample. For example, it may return π_2 shown in Figure 3 obtained by unfolding each loop twice. Then, **Ref** may choose the proof $F_2 = \{\bigvee F \mid F \subseteq \{\phi_0, \phi_1, \phi_2\}\}$ where $\phi_2 \equiv x = a \wedge y = b + 2 \wedge z = 2$ to prove the spuriousness of this new counterexample, which is still insufficient for proving P_{ex} . The abstract-and-refine loop may repeat indefinitely in this manner, adding to the candidates the predicates $F_k = \{\bigvee F \mid F \subseteq \{\phi_i \mid 0 \leq i \leq k\}\}$ where $\phi_i \equiv x = a \wedge y = b + i \wedge z = i$ in each k -th run of the refinement process.

Here, a key observation is that F_{inv} is a proof of safety for *every* counterexample π_1, π_2, \dots of P_{ex} . Because the size of F_{inv} is “small” (under a suitable proof size metric – made more precise below), when using the small refinement heuristic, the refinement process would have to infer F_{inv} (or other small proof of P_{ex}) before producing a large number of incorrect proofs (such as F_1, F_2, \dots).

Our first main result, Theorem 1, states that the above observation holds in general. The result can be proved for a rather generic notion of *proof size*. We

assumed that $size(\cdot)$ satisfies the following: there is a constant $c > 1$ such that for all $n \geq 0$, $|\{F \subseteq \mathcal{T} \mid size(F) \leq n\}| \leq c^n$ (i.e., there are at most c^n proofs of size less than or equal to n). We call such a proof size metric *generic*. The size of the smallest proof for γ , $mpfsize(\gamma)$, is defined to be $\min_{F \in \{F \subseteq \mathcal{T} \mid F \vdash \gamma\}} size(F)$.

Definition 1 (Small Refinement Heuristic). We define $CEGVERIF_{SR}$ to be $CEGVERIF$ with the refinement process Ref satisfying the following property: there is a polynomial f such that for all π and F , if $Ref(\pi)$ returns F then $size(F) \leq f(mpfsize(\pi))$ (i.e., the refinement process returns proofs of size polynomially bounded in that of the smallest proof for the given counterexample).

We are now ready to prove Theorem 1.

Theorem 1. *Let the proof size metric be generic. Then, $CEGVERIF_{SR}$ converges in at most $exp(mpfsize(P))$ many CEGAR iterations given a program P .*

We informally describe the intuition behind the result. First, as remarked above, a proof of a program is also a proof of any of its counterexamples. Therefore, under the small refinement heuristic, the inferred proof in each refinement step is at most polynomially larger than the smallest proof of the program. Then, the result follows from the definition of the generic proof size metric and the fact that the proofs inferred in the refinement process runs must be distinct.

3 Iteration Bound for CFG-Represented Programs

The exponential bound shown in Theorem 1 still seems to have a gap from the performance observed with using the small refinement heuristic in practice. The observation is the motivation for studying a more concrete setting such as CFG-represented programs.

3.1 Preliminaries

Graphs. A finite directed graph $G = (V, E)$ consists of a finite set of nodes V and edges $E \subseteq V \times V$. We write $v(G)$ for V and $e(G)$ for E . For an edge $e = (v, v')$, we write $sc(e)$ for v (the *source* of e) and $tg(e)$ for v' (the *target* of e). Since we only work with finite directed graphs, in what follows, we omit the adjectives and simply write *graphs*.

We write $G \setminus E$ for $(v(G), e(G) \setminus E)$, and $G \setminus V$ for $(V', e(G) \cap V' \times V')$ where $V' = v(G) \setminus V$. For $E \subseteq v(G) \times v(G)$, we write $G \cup E$ for $(v(G), e(G) \cup E)$. A *path* of G is a finite sequence nodes $\varpi = v_1 v_2 \dots v_n$ such that $(v_i, v_{i+1}) \in e(G)$ for each $i \in \{1, \dots, n-1\}$. We write $|\varpi|$ for the length of ϖ and $\varpi(i)$ for the i -th node visited (i.e., $\varpi = \varpi(1)\varpi(2)\dots\varpi(|\varpi|)$).

CFG Programs. We consider programs represented by control flow graphs (also known as *control flow automata* [5]). Formally, a *control flow graph* (CFG) is a tuple (G, T, v_{ini}, v_{err}) where G is a graph, $v_{ini} \in v(G)$ is the *initial node*, $v_{err} \in v(G)$ is the *error node*, and $T : e(G) \rightarrow \mathcal{T}(x, x')$ is the *transition*

relation where $|\mathbf{x}| = |\mathbf{x}'|$. Roughly, each $v \in v(G)$ represents a program location (e.g., a basic block), and for $e \in E$, $T(e)(\mathbf{x}, \mathbf{x}')$ expresses the state transition from the location $sc(e)$ to $tg(e)$ where \mathbf{x} (resp. \mathbf{x}') represents the values of the program variables before (resp. after) the transition. A *path* of the CFG is a path of G . Without loss of generality, we assume that v_{ini} has no incoming edges, v_{err} has no outgoing edges, and all nodes are reachable from v_{ini} . Let \mathbf{Vals} be the set of values. The set of states of the program is $\mathbf{States} = \mathbf{Vals}^{|\mathbf{x}|}$. The set of states reached from $S \subseteq \mathbf{States}$ by taking an edge $e \in e(G)$ is defined to be $Post_T[e](S) = \{s' \in \mathbf{States} \mid s \in S \wedge \models_{\mathcal{T}} T(e)(s, s')\}$. The set of states reached from $S \subseteq \mathbf{States}$ by taking a path ϖ , $Post_T^*[\varpi](S)$, is defined inductively as $Post_T^*[v_1 v_2 \varpi](S) = Post_T^*[v_2 \varpi](Post_T[(v_1, v_2)](S))$ and $Post_T^*[v](S) = Post_T^*[\varepsilon](S) = S$.

We say that a program $P = (G, T, v_{ini}, v_{err})$ is *safe* if for all paths ϖ of P such that $\varpi(1) = v_{ini}$ and $\varpi(|\varpi|) = v_{err}$, $Post_T^*[\varpi](\mathbf{States}) = \emptyset$. In what follows, we often implicitly assume that $|\mathbf{x}, \mathbf{x}'|$ is the arity of the predicates in the range of the transition relation of the CFG being discussed where \mathbf{x} and \mathbf{x}' are distinct variables such that $|\mathbf{x}| = |\mathbf{x}'|$.

We note that the class of CFG programs is already Turing complete when \mathcal{T} is the set of quantifier-free predicates in the theory of linear rational arithmetic (QLRA – *quantifier-free theory of linear rational arithmetic*) (and is equivalent to Boolean programs when \mathcal{T} is propositional), taking the reachable states from \mathbf{States} as the computation result. Checking the safety of CFG programs when $\mathcal{T} = \text{QLRA}$ is undecidable.

Loops. We review the notion of loops in a CFG [2]. A *loop decomposition* of G is a set $\{L_0, \dots, L_n\}$ with each $L_i = (G_i, hds_i)$ satisfying 1.) $G_0 = G$, 2.) each G_i, G_j are either disjoint or one is a subgraph of the other, 3.) each hds_i, hds_j for $i \neq j$ are disjoint, and 4.) each L_i satisfies the following:

- G_i is a non-empty subgraph of G ;
- $G_i \setminus (\{e \in e(G_i) \mid tg(e) \in hds_i\} \cup \bigcup_{j \in Sub(i)} e(G_j)) \cup \bigcup_{j \in Sub(i)} flatten(G_i, G_j)$ is acyclic;
- $hds_i = \{v \in v(G_i) \mid v' \notin v(G_i) \wedge (v', v) \in e(G)\}$; and
- G_i is strongly connected except for G_0

where $flatten(G_i, G_j) = \{(sc(e_1), tg(e_2)) \mid e_1 \in e(G_i) \wedge e_2 \in e(G_j) \wedge tg(e_1) \in v(G_j) \wedge tg(e_2) \in v(G_j)\}$ (i.e., the edges formed by “flattening” G_j in G_i) and $Sub(i) = \{j \mid G_j \text{ is a proper subgraph of } G_i\}$.

Roughly, each loop L_i is a subgraph of G comprising the “back edges” $B_i = \{e \in e(G_i) \mid tg(e) \in hds_i\}$ that take the control flow back to one of the loop entries, and the “loop body” $G_i \setminus B_i$ that is acyclic when the nested subloops are flattened. Note that the loop entries hds_i are the nodes in G_i with incoming edges from nodes outside of the loop. A loop decomposition forms a tree with L_0 as the root and nested subloops as children.³

³ In the literature, L_0 is typically not treated as a loop, and a loop decomposition forms a forest.

A graph is *rooted* if there is a node with no incoming edges and from which every node in the graph is reachable. Clearly, the graph underlying a CFG is rooted. The following is a folk theorem (cf. Appendix A.4 for a proof).

Theorem 4. *A rooted graph has a loop decomposition.*

We define a loop decomposition of a CFG (G, T, v_{ini}, v_{err}) to be a loop decomposition of G . In the following, we assume that each CFG P is associated with its loop decomposition $loops(P)$ (e.g., constructed by the algorithm given in Appendix A.4). We write $maxhds(P)$ for $\max_{(.,hds) \in loops(P)} |hds|$. We remark that $maxhds(P) \leq 1$ if P is reducible [2].

Proofs and Counterexamples for CFG Programs. Informally, a counterexample of a CFG is an acyclic CFG obtained by *unfolding* the loops of the CFG. To formalize loop unfolding, we introduce a simple graph grammar below.

$$\begin{aligned} instr &::= v_0 \mapsto (\phi_1 : v_1, \phi_2 : v_2, \dots, \phi_n : v_n) \\ t &::= \emptyset \mid \{instr\} \mid t_1 \cup t_2 \mid loop(h_1, \dots, h_m) t \end{aligned}$$

Here, in each *instruction* $v_0 \mapsto (\phi_1 : v_1, \dots, \phi_n : v_n)$, v_1, \dots, v_n are distinct and $\phi_i \in \mathcal{T}(\mathbf{x}, \mathbf{x}')$ for each $i \in \{1, \dots, n\}$. We call v_0 the *source* and v_1, \dots, v_n the *targets* of the instruction. Roughly, $v_0 \mapsto (\phi_1 : v_1, \dots, \phi_n : v_n)$ expresses the set of edges $\{(v_0, v_1), \dots, (v_0, v_n)\}$ such that the transition relation of (v_0, v_i) is ϕ_i . A *term* t expresses the CFG comprising the edges represented by the instructions occurring in t . The *loop* annotations mark subparts of the CFG that correspond to loops, so that $loop(\mathbf{h}) t$ expresses a loop with \mathbf{h} being the entry nodes and t representing the union of the loop body, the back edges, and the edges from the loop body to the outer loops. We require the sources of the instructions occurring in a term to be distinct. We refer to Appendix B for the formal correspondence between CFGs and graph grammar terms. In what follows, we equate a CFG with the corresponding graph grammar term.

We overload t for the set of instructions occurring in t . We let $sc(v_0 \mapsto (\phi_1 : v_1, \dots, \phi_n : v_n)) = \{v_0\}$, and $tg(v_0 \mapsto (\phi_1 : v_1, \dots, \phi_n : v_n)) = \{v_1, \dots, v_n\}$. We let $sc(t) = \bigcup_{instr \in t} sc(instr)$, and $tg(t) = \bigcup_{instr \in t} tg(instr)$. We let $v(t) = sc(t) \cup tg(t)$. For sets of nodes V_1 and V_2 , we write $t(V_1, V_2)$ for t with each source occurrence of $v \in V_1$ replaced by ' v ' and each target occurrence of $v \in V_2$ replaced by ' v '. For instance, for $t = \{v_0 \mapsto (\top : v_1, \top : v_2), v_1 \mapsto (\top : v_0, \top : v_3)\}$, $t(\{v_0\}, \{v_1, v_3\}) = \{v_0 \mapsto (\top : v_1, \top : v_2), v_1 \mapsto (\top : v_0, \top : v_3)\}$. We write $t \setminus V$ for t with each instruction $v_0 \mapsto (v_1 : \phi_1, \dots, v_n : \phi_n)$ replaced by $v_0 \mapsto (v_{i1} : \phi_{i1}, \dots, v_{im} : \phi_{im})$ where $\{v_{i1}, \dots, v_{im}\} = \{v_1, \dots, v_n\} \setminus V$, treating an instruction with empty targets as \emptyset . For $\mathbf{h} = h_1, \dots, h_n$, we write ' \mathbf{h} ' for ' h_1, \dots, h_n '.

Let rewriting contexts be defined as follows.

$$C ::= [] \mid C \cup t \mid loop(h_1, \dots, h_m) C$$

Loop unfolding is defined by the following rewriting rules.

$$\begin{aligned} C[loop(\mathbf{h}) t] &\dashrightarrow C[t(v(t) \setminus \{\mathbf{h}\}, sc(t)) \cup loop(\mathbf{h}) t \setminus \{\mathbf{h}\}, \{\mathbf{h}\}] \\ C[loop(\mathbf{h}) t] &\dashrightarrow C[t \setminus \{\mathbf{h}\}] \end{aligned}$$

The first rule unfolds the loop once, whereas the second rule “closes” the loop by removing the back edges. We formalize the *counterexamples* of a CFG P , $\text{cex}(P)$, to be the acyclic CFGs obtained by applying the above rewriting an arbitrary number of times, that is, $\text{cex}(P) = \{\pi \mid P \dashrightarrow^* \pi \wedge \pi \text{ is acyclic}\}$.

A proof of a CFG program or counterexample is a set of predicates that forms a Floyd-style node-wise inductive invariant [9]. More formally, we say that $\sigma : v(G) \rightarrow \mathcal{T}(\mathbf{x})$ is a *node-wise inductive invariant* of $\gamma = (G, T, v_{ini}, v_{err})$, written $\sigma \vdash_{cf\!g} \gamma$, if 1.) $\sigma(v_{ini}) = \top$, 2.) $\sigma(v_{err}) = \perp$, and 3.) for each $e \in e(G)$, $\models_{\mathcal{T}} \sigma(sc(e))(\mathbf{x}) \wedge T(e)(\mathbf{x}, \mathbf{x}') \Rightarrow \sigma(tg(e))(\mathbf{x}')$. We say that F is a *proof* of γ , written $F \vdash_{cf\!g} \gamma$, if there exists $\sigma : v(G) \rightarrow F \cup \{\perp, \top\}$ such that $\sigma \vdash_{cf\!g} \gamma$. The following are immediate from the definition of $\vdash_{cf\!g}$.

Theorem 5 (Soundness of $\vdash_{cf\!g}$). *If $F \vdash_{cf\!g} \gamma$ then γ is safe.*

Theorem 6 (Monotonicity of $\vdash_{cf\!g}$). *If $F \vdash_{cf\!g} \gamma$ and $F \subseteq F'$, then $F' \vdash_{cf\!g} \gamma$.*

Theorem 7. *Suppose $F \vdash_{cf\!g} P$. Then, for any $\pi \in \text{cex}(P)$, $F \vdash_{cf\!g} \pi$.*

Theorems 5, 6 and 7 justify us to use $\vdash_{cf\!g}$ for the proof relation \vdash of CEGVERIF (cf. Section 2.1).

Next, we concretize the abstraction process and counterexample generation for CFGs. We present the Cartesian predicate abstraction, a form of predicate abstraction used in CEGAR, as the abstraction process. Then, we present a counterexample set that is *sound* for the abstraction process to generate and refute (made precise below). We note that, because our result is shown for the Cartesian predicate abstraction, it also holds for stronger abstraction processes such as the Boolean predicate abstraction.

We write F^\wedge for the \wedge -closure of F (i.e., $\{\bigwedge F' \mid F' \subseteq F\}$) and $F^{\wedge\vee}$ for the $\wedge\vee$ -closure of F (i.e., $\{\bigvee F' \mid F' \subseteq F^\wedge\}$). We write F_\perp for $F \cup \{\perp\}$. For $\gamma = (G, T, v_{ini}, v_{err})$, $\sigma : v(G) \rightarrow \mathcal{T}$, and $F \subseteq \mathcal{T}$, we say that σ is a *F -Cartesian predicate abstraction node-wise inductive invariant* of γ , written $\sigma \vdash_{crt}^F \gamma$, if 1.) $\sigma(v_{ini}) = \top$, 2.) $\sigma(v_{err}) = \perp$, and 3.) for each $e \in e(G)$, we have $\sigma(sc(e)) = \bigvee F_1$, $\sigma(tg(e)) = \bigvee F_2$ for some $F_1 \subseteq F_\perp^\wedge$ and $F_2 \subseteq F_\perp^\wedge$ such that for each $\phi \in F_1$, there is $\psi \in F_2$ where $\models_{\mathcal{T}} \phi(\mathbf{x}) \wedge T(e)(\mathbf{x}, \mathbf{x}') \Rightarrow \psi(\mathbf{x}')$. We write $F \vdash_{crt} \gamma$ if there exists σ such that $\sigma \vdash_{crt}^F \gamma$. Clearly, $F \vdash_{cf\!g} \gamma$ implies $F \vdash_{crt} \gamma$, and $F \vdash_{crt} \gamma$ implies $F^{\wedge\vee} \vdash_{cf\!g} \gamma$.

The *Cartesian predicate abstraction* abstraction process $\text{Abs}^{crt}(\cdot, F)$ is an abstract interpretation [7] over the finite lattice $(F^{\wedge\vee}, \models_{\mathcal{T}} \cdot \Rightarrow \cdot)$ with the abstract state transformer $\alpha^F(Post)$ that computes the strongest cube over F_\perp implied in the next state (from the initial abstract state \top):

$$\alpha^F(Post)_T[e](\phi) = \bigwedge \{\psi \in F_\perp \mid \models_{\mathcal{T}} \phi(\mathbf{x}) \wedge T(e)(\mathbf{x}, \mathbf{x}') \Rightarrow \psi(\mathbf{x}')\}$$

The abstraction process guarantees 1.) $F \vdash_{crt} \gamma$ if and only if $\text{Abs}^{crt}(\gamma, F)$ returns safe, and 2.) if $\text{Abs}^{crt}(P, F)$ returns a counterexample π , then $F \not\vdash_{crt} \pi$.

We show that Abs^{crt} satisfies the requirements of the abstraction process given in Section 2.1.

Theorem 8. Abs^{crt} satisfies the requirements for Abs. That is,

- If $F \vdash_{cfg} \gamma$ then $\text{Abs}^{crt}(\gamma, F) = \text{safe}$;
- If $\text{Abs}^{crt}(\gamma, F)$ returns π then $\text{Abs}^{crt}(\pi, F) \neq \text{safe}$;
- If $F \vdash_{cfg} P$ and $\text{Abs}^{crt}(P, F')$ returns π then $F \vdash_{cfg} \pi$; and
- If $\text{Abs}^{crt}(\gamma, F) = \text{safe}$ then $\exists F' \subseteq \mathcal{T}. F' \vdash_{cfg} \gamma$.

We say that the set of counterexamples $\mathcal{X}(P) \subseteq \text{cex}(P)$ is *sound* for Abs if $\text{Abs}(P, F) \neq \text{safe}$ implies that there exists $\pi \in \mathcal{X}(P)$ such that $F \not\vdash_{cfg} \pi$ (i.e., generating and refuting only the counterexamples from $\mathcal{X}(P)$ is sufficient for verifying P). We say \mathcal{X} is sound for Abs if $\mathcal{X}(P)$ is sound for Abs for each P . Let $\text{cex}^{syn}(P) \subseteq \text{cex}(P)$ be the set of counterexamples obtained by unfolding the loops $\text{loops}(P)$ an arbitrary but the same number of times for each loop (copies of nested subloops are also unfolded the same number of times – cf. Appendix C for the formal definition). We show that cex^{syn} is sound for Abs^{crt} .

Theorem 9. cex^{syn} is sound for Abs^{crt} .

Theorems 8 and 9 justify us to use Abs^{crt} with cex^{syn} as the counterexample generator for the abstraction process of CEGVERIF (cf. Section 2.1).

We note that, in the setting described above, a counterexample can be a general dag-shaped CFG. While early CEGAR verifiers often restricted the counterexamples to paths [4, 11], more recent verifiers also use dag counterexamples, and researchers have proposed methods for inferring small refinements from such counterexamples [8, 10, 19, 16, 22]. Also, we note that, when \mathcal{T} is QFLRA, checking the provability for a dag CFG γ (i.e., checking if $\exists F \subseteq \mathcal{T}. F \vdash_{cfg} \gamma$) is decidable, whereas it is undecidable for an arbitrary (i.e., cyclic) CFG [14].⁴

3.2 Main Result

We define the *syntactic size* of F to be the sum of the syntactic sizes of the predicates, that is, $\text{size}(F) = \sum_{\phi \in F} |\phi|$ where $|\phi|$ is the number of (logical and non-logical) symbols in ϕ . We state the main result.

Theorem 10. Let the proof size metric be syntactic, \vdash_{cfg} be the proof relation, and Abs^{crt} be the abstraction process with cex^{syn} as the counterexample generator. Then, CEGVERIF_{SR} converges in $\text{poly}(\text{mpfsize}(P))^{\text{maxhds}(P)}$ many CEGAR iterations given a CFG program P .

We informally describe the intuition behind the result by analyzing the behavior of CEGVERIF_{SR} on the program P_{ex} from Figure 2, treated as a CFG program. Figure 4 (a) shows the CFG representation of the program. Here, node 1 is the initial node, node 4 is the error node, and the transition relation is as follows.

$$\begin{aligned}
\phi_1(\mathbf{x}, \mathbf{x}') &\equiv x' = a' \wedge y' = b' \wedge z' = 0 \\
\phi_2(\mathbf{x}, \mathbf{x}') &\equiv z' = z + 1 \wedge y' = y + 1 \wedge x' = x \wedge a' = a \wedge b' = b \\
\phi_3(\mathbf{x}, \mathbf{x}') &\equiv x' = x \wedge y' = y \wedge z' = z \wedge a' = a \wedge b' = b \\
\phi_4(\mathbf{x}, \mathbf{x}') &\equiv z \neq 0 \wedge z' = z - 1 \wedge y' = y - 1 \wedge x' = x \wedge a' = a \wedge b' = b \\
\phi_5(\mathbf{x}, \mathbf{x}') &\equiv z = 0 \wedge a = b \wedge x \neq y
\end{aligned}$$

⁴ The decidability holds for any theory with effective interpolation (cf. Appendix D).




Fig. 4. (a) CFG representation of P_{ex} and (b) CFG representation of $\text{cex}^{\text{syn}}(P_{ex})$

where $\mathbf{x} = x, y, z, a, b$ and $\mathbf{x}' = x', y', z', a', b'$. P_{ex} has two non-root loops L_1 and L_2 such that L_1 corresponds to the first loop (lines 3-5) and L_2 corresponds to the second loop (lines 6-8) of Figure 2. An element of $\text{cex}^{\text{syn}}(P_{ex})$ is a CFG of the form shown in Figure 4 (b) where each loop is unfolded k times for some $k > 0$. Nodes $2_1, \dots, 2_k$ (resp. $3_1, \dots, 3_k$) are the copies of the entry node of L_1 (resp. L_2) created by the unfolding.

Recall that the small refinement heuristic returns a proof of size polynomially bounded in the size of the smallest proof of the given counterexample. Let f be the polynomial factor (i.e., Ref returns proofs of size at most $f(\text{mpfsize}(\pi))$ given a counterexample π). Then, because F_{inv} is a proof of any counterexample of P_{ex} , a proof returned by Ref in a run of $\text{CEGVERIF}_{\text{SR}}(P_{ex})$ would be of size at most $f(\text{size}(F_{\text{inv}}))$. Consider the counterexample $\pi_n \in \text{cex}^{\text{syn}}(P_{ex})$ obtained by unfolding each loop n times where $n > f(\text{size}(F_{\text{inv}}))$. It can be seen that any proof F of π_n such that $\text{size}(F) \leq f(\text{size}(F_{\text{inv}}))$ must form a Cartesian predicate abstraction node-wise inductive invariant of P_{ex} (i.e., $F \vdash_{\text{crt}} P_{ex}$). This is because for such a F to not be a \vdash_{crt} -proof of P_{ex} , it must assign distinct predicates to the unfolded loop heads of each loop, but that is not possible when the proof size must be at most $f(\text{size}(F_{\text{inv}}))$.

More generally, as in Theorem 1, the first key observation used to prove Theorem 10 is the fact that a proof for the program is a proof for any counterexample of the program. As before, this implies that the small refinement heuristic ensures that the proofs inferred in the refinement are only polynomially larger than the smallest proof for the program. Then, we use the observation

that, for a CFG program, if a counterexample is “sufficiently large”, then *any* small proof for the counterexample becomes a Cartesian predicate abstraction node-wise inductive invariant for the program, and so CEGAR will have to halt in the next iteration. Finally, the iteration bound follows from the number of counterexamples in cex^{syn} that are within such a sufficiently large size.

We remark that it is actually sufficient for the proof size metric to only satisfy $\text{size}(F) \geq |F|$ for Theorem 10 to hold (i.e., at least the number of predicates in F). Also, as remarked before, the result also holds for stronger abstraction processes such as the Boolean predicate abstraction. Meanwhile, as we shall further discuss in Section 4, a limitation of the result is that it only considers counterexamples where all loops are unfolded the same number of times. This implies that the abstraction process may generate large counterexamples in relatively early stages of the verification, especially when the program contains nested loops (cf. Appendix C).

We end the section by showing that the proof size metric is crucial to the result. That is, under the generic proof size metric (cf. Section 2.2), we can only guarantee an exponential bound (which is ensured by the generic result of Theorem 1), even when the rest of the setting is left unchanged from Theorem 10.

Theorem 11. *Let the proof size metric be generic, \vdash_{cfg} be the proof relation, and Abs^{crt} be the abstraction process with cex^{syn} as the counterexample generator. Then, there exists a CFG program P with a proof $F \vdash_{\text{cfg}} P$ on which $\text{CEGVERIF}_{\text{SR}}$ may take $\exp(\text{size}(F))$ many CEGAR iterations to converge.*

4 Limitations

While we believe that the paper’s results are a step toward understanding the effectiveness of the small refinement heuristic, we still have ways to go to get the whole picture. Below, we discuss some limitations of our work.

First, we do not account for the cost of the abstraction process and the refinement process (i.e., we only show bounds on the number of CEGAR iterations). For instance, the running time of the abstraction process typically grows as the number of candidate predicates grows. Also, previous work has suggested that inferring a small proof of a counterexample may be computationally expensive (cf. Appendix F of [20]), and while there has been much progress on efficient algorithms for inferring a small proof of a counterexample [12, 3, 17, 22], explaining why such algorithms work well in practice seems to be no easier than explaining the efficiency of the overall verification process.⁵

Secondly, in our setting of CFG programs, the counterexample form is rather restricted. That is, we only consider counterexamples obtained by unfolding the

⁵ In a sense, this paper poses and studies the question “assuming we have such algorithms for inferring small refinements, what can be said about the overall verification efficiency?”. Note that a possible outcome of the study can be a negative result; for example, showing that inferring small refinements is hard because otherwise it would give an efficient algorithm to some provably hard verification problem.

program’s loops, and unfolding only the same number of times for every (copy of) loops (i.e., cex^{syn}). By contrast, an actual CEGAR verifier is often more liberal about the counterexample forms, and for example, may allow arbitrary unfoldings or paths as counterexamples.⁶

5 Conclusion

We have presented a theoretical explanation for the effectiveness of the small refinement heuristic used in program verification with CEGAR, which, to our knowledge, is the first of its kind. Specifically, we have formalized the small refinement heuristic to be a refinement process that returns a proof whose size is polynomially bounded in that of the smallest proof for the given counterexample, and shown that CEGAR with such a refinement is guaranteed to converge in a number of iterations bounded by the size of a proof of the given program. We have presented the results under a rather generic setting of CEGAR, and under a more concrete setting for CFG-represented programs.

Acknowledgements. We thank the anonymous reviewers for useful comments. This work was supported by MEXT Kakenhi 26330082 and 25280023, and JSPS Core-to-Core Program, A.Advanced Research Networks.

References

1. International competition on software verification (SV-COMP). <http://sv-comp.sosy-lab.org/>.
2. A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1986.
3. A. Albarghouthi and K. L. McMillan. Beautiful interpolants. In N. Sharygina and H. Veith, editors, *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 313–329. Springer, 2013.
4. T. Ball and S. K. Rajamani. The SLAM project: debugging system software via static analysis. In J. Launchbury and J. C. Mitchell, editors, *POPL*, pages 1–3. ACM, 2002.
5. D. Beyer, A. Cimatti, A. Griggio, M. E. Keremoglu, and R. Sebastiani. Software model checking via large-block encoding. In *FMCAD*, pages 25–32. IEEE, 2009.
6. S. A. Cook. The complexity of theorem-proving procedures. In M. A. Harrison, R. B. Banerji, and J. D. Ullman, editors, *STOC*, pages 151–158. ACM, 1971.
7. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In R. M. Graham, M. A. Harrison, and R. Sethi, editors, *POPL*, pages 238–252. ACM, 1977.
8. J. Esparza, S. Kiefer, and S. Schwoon. Abstraction refinement with Craig interpolation and symbolic pushdown systems. *JSAT*, 5(1-4):27–56, 2008.

⁶ It is easy to reduce any CFG program to an equivalent one whose unfoldings/paths would coincide with cex^{syn} (e.g., by encoding program locations in transition relation – cf. Appendix E). But, such a reduction is likely to affect the cost of the abstraction process and the refinement process.

9. R. W. Floyd. Assigning meanings to programs. In *Symposia in Applied Mathematics*, volume 19, pages 19–32, 1967.
10. B. S. Gulavani, S. Chakraborty, A. V. Nori, and S. K. Rajamani. Refining abstract interpretations. *Information Processing Letters*, 110(16):666–671, 2010.
11. T. A. Henzinger, R. Jhala, R. Majumdar, and K. L. McMillan. Abstractions from proofs. In N. D. Jones and X. Leroy, editors, *POPL*, pages 232–244. ACM, 2004.
12. K. Hoder, L. Kovács, and A. Voronkov. Playing in the grey area of proofs. In J. Field and M. Hicks, editors, *POPL*, pages 259–272. ACM, 2012.
13. R. Jhala and K. L. McMillan. A practical and complete approach to predicate refinement. In H. Hermanns and J. Palsberg, editors, *TACAS*, volume 3920 of *Lecture Notes in Computer Science*, pages 459–473. Springer, 2006.
14. N. Kobayashi. Personal communication, 2012.
15. K. L. McMillan. Quantified invariant generation using an interpolating saturation prover. In C. R. Ramakrishnan and J. Rehof, editors, *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 413–427. Springer, 2008.
16. P. Rümmer, H. Hojjat, and V. Kuncak. Disjunctive interpolants for Horn-clause verification. In N. Sharygina and H. Veith, editors, *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 347–363. Springer, 2013.
17. C. Scholl, F. Pigorsch, S. Disch, and E. Althaus. Simple interpolants for linear arithmetic. In *DATE*, pages 1–6. IEEE, 2014.
18. R. Tarjan. Testing flow graph reducibility. In *STOC*, pages 96–107, 1973.
19. T. Terauchi. Dependent types from counterexamples. In M. V. Hermenegildo and J. Palsberg, editors, *POPL*, pages 119–130. ACM, 2010.
20. T. Terauchi and H. Unno. Relaxed stratification: A new approach to practical complete predicate refinement. In J. Vitek, editor, *ESOP*, volume 9032 of *Lecture Notes in Computer Science*, pages 610–633. Springer, 2015.
21. H. Unno and N. Kobayashi. Dependent type inference with interpolants. In A. Porto and F. J. López-Fraguas, editors, *PPDP*, pages 277–288. ACM, 2009.
22. H. Unno and T. Terauchi. Inferring simple solutions to recursion-free Horn clauses via sampling. In C. Baier and C. Tinelli, editors, *TACAS*, volume 9035 of *Lecture Notes in Computer Science*, pages 149–163. Springer, 2015.

A Proofs

A.1 Proof of Theorem 2

Theorem 2. $\text{CEGVERIF}(P)$ returns safe only if $\exists F \subseteq \mathcal{T}.F \vdash P$.

Proof. Immediate from the fact that $\text{Abs}(P, F)$ returns safe only if $\exists F \subseteq \mathcal{T}.F \vdash P$. \square

A.2 Proof of Theorem 3

Theorem 3. $\text{CEGVERIF}(P)$ returns unsafe only if $\forall F \subseteq \mathcal{T}.F \not\vdash P$.

Proof. We show the contrapositive. Suppose $F \vdash P$. Then $F \vdash \pi$ for any π returned by $\text{Abs}(F', P)$ for some F' . Hence, for any counterexample π passed to Ref in the run of $\text{CEGVERIF}(P)$, $F \vdash \pi$. Therefore, by the property of Ref that it returns unsafe only if given a counterexample π such that $\forall F \subseteq \mathcal{T}.F \not\vdash \pi$, it follows that $\text{CEGVERIF}(P)$ does not return unsafe. \square

A.3 Proof of Theorem 1

The following lemma states that each run of the refinement process returns a new proof.

Lemma 1 (Progress). *Suppose $\text{Abs}(P, F)$ returns π and $\text{Ref}(\pi)$ returns F' . Then, $F' \not\subseteq F$.*

Proof. Because $\text{Ref}(\pi)$ returns F' , we have $F' \vdash \pi$. Also, because $\text{Abs}(P, F)$ returns π , we have $\text{Abs}(\pi, F) \neq \text{safe}$, and so $F \not\vdash \pi$. Therefore, $F \neq F'$, and by the monotonicity of \vdash , it follows that $F' \not\subseteq F$. \square

Theorem 1. *Let the proof size metric be generic. Then, $\text{CEGVERIF}_{\text{SR}}$ converges in at most $\exp(\text{mpfsize}(P))$ many CEGAR iterations given a program P .*

Proof. Let F_P be a proof of P such that $\text{mpfsize}(P) = \text{size}(F_P)$. Then, for any counterexample π of P (i.e., $\text{Abs}(P, F')$ returns π for some F'), $F_P \vdash \pi$, and therefore $\text{mpfsize}(\pi) \leq \text{size}(F_P)$. There are at most $c^{f(\text{mpfsize}(P))}$ proofs of size at most $f(\text{size}(F_P))$. Therefore, by Lemma 1, $\text{CEGVERIF}_{\text{SR}}$ converges in at most $\exp(\text{mpfsize}(P))$ many CEGAR iterations. \square

A.4 Proof of Theorem 4

Algorithm 1 Loop Decomposition Construction

Input: Rooted graph G

Output: Loop Decomposition of G

```

1:  $decomp := \{(G, \emptyset)\}$ ;
2:  $\text{PROCESS}(G)$ ;
3: return  $decomp$ 
4: procedure  $\text{PROCESS}(G')$ 
5:    $\{G_1, \dots, G_m\} = \text{SCC}(G')$ ;
6:   for each cyclic  $G_i$  do
7:      $hds = \{v \in v(G_i) \mid v' \notin v(G_i) \wedge (v', v) \in e(G)\}$ ;
8:      $decomp := decomp \cup \{(G_i, hds)\}$ ;
9:      $\text{PROCESS}(G_i \setminus hds)$ 
10:  end for
11: end procedure

```

Theorem 4. *A rooted graph has a loop decomposition.*

The result is well known for reducible graphs [18].⁷ Below, we provide a proof for the general (i.e., possibly irreducible) case.

⁷ For a reducible graph, each loop has only one entry node.

Proof. Algorithm 1 builds the loop decomposition *decomp* by first initializing it to contain just $L_0 = (G, \emptyset)$, and then calling the recursive procedure PROCESS on the given graph. In PROCESS, the helper procedure SCC returns the set of strongly connected components of the given graph.

Because G is rooted, *hds* in the **for each** block of PROCESS is always non-empty, and therefore, the recursive call to PROCESS is on a smaller graph, and the algorithm terminates. Then, by induction on the recursive call depth, it follows that the constructed *decomp* is a loop decomposition of G . \square

A.5 Proof of Theorem 5

Theorem 5. *If $F \vdash_{cfg} \gamma$ then γ is safe.*

Proof. Let $\gamma = (G, T, v_{ini}, v_{err})$. Suppose $\sigma \vdash_{cfg} \gamma$ where $\sigma : v(G) \rightarrow F \cup \{\perp, \top\}$. Let ϖ be a path of γ such that $\varpi(1) = v_{ini}$ and $\varpi(|\varpi|) = v_{err}$. Let $\sigma_{\varpi} : \{1, \dots, |\varpi|\} \rightarrow F \cup \{\perp, \top\}$ be such that $\sigma_{\varpi}(i) = \sigma(\varpi(i))$ for each $i \in \{1, \dots, |\varpi|\}$. Then, for any prefix ϖ' of ϖ , by induction on the length of ϖ' , we have $Post_{\mathcal{T}}^*[\varpi'](\text{States}) \subseteq \{s \in \text{States} \mid \models_{\mathcal{T}} \sigma_{\varpi}(|\varpi'|)(s)\}$. Therefore, $Post_{\mathcal{T}}^*[\varpi](\text{States}) \subseteq \{s \in \text{States} \mid \models_{\mathcal{T}} \sigma_{\varpi}(|\varpi|)(s)\} = \emptyset$. \square

A.6 Proof of Theorem 6

Theorem 6. *If $F \vdash_{cfg} \gamma$ and $F \subseteq F'$, then $F' \vdash_{cfg} \gamma$.*

Proof. Immediate from the definition of \vdash_{cfg} . \square

A.7 Proof of Theorem 7

Theorem 7. *Suppose $F \vdash_{cfg} P$. Then, for any $\pi \in \text{cex}(P)$, $F \vdash_{cfg} \pi$.*

Proof. Let $P = (G, T, v_{ini}, v_{err})$. Let $\pi \in \text{cex}(P)$ where $\pi = (G_{\pi}, T_{\pi}, v_{ini}, v_{err})$. Suppose $\sigma \vdash_{cfg} P$ where $\sigma : v(G) \rightarrow F \cup \{\perp, \top\}$. Let $\sigma_{\pi} : v(G_{\pi}) \rightarrow F \cup \{\perp, \top\}$ such that $\sigma_{\pi}(v') = \sigma(v)$ where $v' \in v(G_{\pi})$ is the copy of the node $v \in v(G)$ created by the unfolding. Then, from the construction of π , we have $\sigma_{\pi} \vdash_{cfg} \pi$. \square

A.8 Proof of Theorem 8

Theorem 8. *Abs^{crt} satisfies the requirements for Abs. That is,*

- (1) *If $F \vdash_{cfg} \gamma$ then $\text{Abs}^{crt}(\gamma, F) = \text{safe}$;*
- (2) *If $\text{Abs}^{crt}(\gamma, F)$ returns π then $\text{Abs}^{crt}(\pi, F) \neq \text{safe}$;*
- (3) *If $F \vdash_{cfg} P$ and $\text{Abs}^{crt}(P, F')$ returns π then $F \vdash_{cfg} \pi$; and*
- (4) *If $\text{Abs}^{crt}(\gamma, F) = \text{safe}$ then $\exists F' \subseteq \mathcal{T}. F' \vdash_{cfg} \gamma$.*

Proof. We show each (1)-(4).

- (1) Immediate from the fact that $F \vdash_{cfg} \gamma$ implies $F \vdash_{crt} \gamma$.
- (2) Suppose $\text{Abs}^{crt}(\gamma, F)$ returns π . Then, $F \not\vdash_{crt} \pi$. Hence, we have $\text{Abs}^{crt}(\pi, F) \neq \text{safe}$.
- (3) Suppose $F \vdash_{cfg} P$ and $\text{Abs}^{crt}(P, F')$ returns π . Then, $\pi \in \text{cex}(P)$, and by Theorem 7, we have $F \vdash_{cfg} \pi$.
- (4) Immediate from the fact that $F \vdash_{crt} \gamma$ implies $F^{\wedge\vee} \vdash_{cfg} \gamma$.

□

A.9 Proof of Theorem 9

Theorem 9. cex^{syn} is sound for Abs^{crt} .

Proof. Suppose that $F \vdash_{cfg} \pi$ for all $\pi \in \text{cex}^{syn}(P)$. It suffices to show that $\text{Abs}(P, F) = \text{safe}$.

Then, for some “sufficiently large” $\pi \in \text{cex}^{syn}(P)$, we have $F \vdash_{cfg} \pi$. Therefore, by the construction of a Cartesian predicate abstraction node-wise inductive invariant from a proof of a sufficiently large counterexample (cf. Appendix A.10), we have $F \vdash_{crt} P$, and so $\text{Abs}^{crt}(P, F) = \text{safe}$. □

A.10 Proof of Theorem 10

Theorem 10. Let the proof size metric be syntactic, \vdash_{cfg} be the proof relation, and Abs^{crt} be the abstraction process with cex^{syn} as the counterexample generator. Then, $\text{CEGVERIF}_{\text{SR}}$ converges in $\text{poly}(\text{mpfsize}(P))^{\text{maxhds}(P)}$ many CEGAR iterations given a CFG program P .

Proof. Let F_P be a proof of $P = (G, T, v_{ini}, v_{err})$ such that $\text{mpfsize}(P) = \text{size}(F_P)$. Then, for any counterexample π of P (i.e., $\text{Abs}^{crt}(P, F)$ returns π for some F), $F_P \vdash \pi$, and so $\text{mpfsize}(\pi) \leq \text{size}(F_P)$. Let $\text{lim} = (f(\text{size}(F_P)) + 2)^{\text{maxhds}(P)} + 1$ (“+2” accounts for $\{\perp, \top\}$). Let π_{lim} be a counterexample in $\text{cex}^{syn}(P)$ obtained by unfolding the loops at least lim many times. We show that for any F such that $F \vdash_{cfg} \pi_{\text{lim}}$ and $\text{size}(F) \leq f(\text{size}(F_P))$, $F \vdash_{crt} P$. (We call such a counterexample π_{lim} *sufficiently large*.) Then, the result follows from the fact that there are only lim many counterexamples in $\text{cex}^{syn}(P)$ that are obtained by unfolding the loops at most lim many times, and the fact that no counterexample is returned more than once by the refinement process in a run of $\text{CEGVERIF}_{\text{SR}}$.

Let $\pi_{\text{lim}} = (G_{\text{lim}}, T_{\text{lim}}, v_{ini}, v_{err})$. Let $\sigma_{\text{lim}} : v(G_{\text{lim}}) \rightarrow F \cup \{\perp, \top\}$ be such that $\text{size}(F) \leq f(\text{size}(F_P))$ and $\sigma_{\text{lim}} \vdash_{cfg} \pi_{\text{lim}}$. Let $k \geq \text{lim}$ be the number of times the loops are unfolded in π_{lim} . We construct $\sigma : v(G) \rightarrow F^{\wedge\vee}$ such that $\sigma \vdash_{crt}^F P$ by “folding” the unfolded loops in a bottom up manner. We initialize $\sigma = \sigma_{\text{lim}}$, and $\gamma = \pi_{\text{lim}}$. We iteratively fold γ from leaf loops, while maintaining the property that $\sigma \vdash_{crt}^F \gamma$. Then, the result follows from the fact that γ becomes P at the root of the folding process.

Let γ and σ be the current CFG and its Cartesian predicate abstraction node-wise inductive invariant (i.e., $\sigma \vdash_{crt}^F \gamma$). Let $\gamma = C[t'_1 \cup t'_2 \cup \dots \cup t'_k \cup t_k \setminus \{\mathbf{h}_k\}]$ where

$\mathbf{h}_0 = \mathbf{h}$, $t_0 = t$, $\mathbf{h}_{i+1} = \langle \mathbf{h}_i, t_i' \rangle$, $t_i' = t_i \langle v(t_i) \setminus \{\mathbf{h}_i\}, sc(t_i) \rangle$, and $t_{i+1} = t_i \langle \mathbf{h}_i, \mathbf{h}_i \rangle$ for each $i \in \{0, \dots, k-1\}$. Let $\gamma' = C[\text{loop } (\mathbf{h}) t]$. That is, γ' is obtained by folding the unfolded loop of γ . We construct $\sigma' : v(\gamma') \rightarrow F^{\wedge V}$ such that $\sigma' \vdash_{crt}^F \gamma'$ as follows. Because $\text{ran}(\sigma_{lim}) = F \cup \{\perp, \top\}$, for some $1 \leq i_1 < i_2 \leq (f(\text{size}(F_P)) + 2)^{|\mathbf{h}|} + 1 \leq k$, $\sigma_{lim}(h^{i_1}) = \sigma_{lim}(h^{i_2})$ for each $h \in \{\mathbf{h}\}$. By construction, $\sigma_{lim}(h^i) = \sigma(h^i)$ for each $i \in \{1, \dots, k\}$ and $h \in \{\mathbf{h}\}$. We set $\sigma'(v) = \bigvee_{i=1}^{i_2} \sigma(v^i)$ for each $v \in sc(t)$, and $\sigma'(v) = \sigma(v)$ for each $v \in v(\gamma) \setminus sc(t)$. Then, it can be seen that $\sigma' \vdash_{crt}^F \gamma'$. \square

A.11 Proof of Theorem 11

Theorem 11. *Let the proof size metric be generic, \vdash_{cfg} be the proof relation, and Abs^{crt} be the abstraction process with cex^{syn} as the counterexample generator. Then, there exists a CFG program P with a proof $F \vdash_{cfg} P$ on which $\text{CEGVERIF}_{\text{SR}}$ may take $\text{exp}(\text{size}(F))$ many CEGAR iterations to converge.*

Proof. We show that P_{ex} from Figure 4 is such a program. As remarked before, we have $F_{inv} \vdash_{cfg} P_{ex}$ where $F_{inv} = \{a = b \Rightarrow y = x + z\}$. For each $k > 0$, let $\pi_k \in \text{cex}^{syn}(P_{ex})$ be the counterexample obtained by unfolding each loop k times, and let $F_k = \{\bigvee F \mid F \subseteq \{\phi_i \mid 0 \leq i \leq k\}\}$ where $\phi_i \equiv x = a \wedge y = b + i \wedge z = i$. Then, $F_k \vdash_{cfg} \pi_k$ for each $k > 0$, and $\bigcup_{i=1}^j F_i \not\vdash_{crt} \pi_k$ for each $k > j > 0$.

Let $\ell > 0$. Define size as follows: $\text{size}(F_{inv}) = \ell$, $\text{size}(F_k) = \lfloor \log k \rfloor$ for $k \in \{1, \dots, 2^\ell - 1\}$, and $\text{size}(F') = \ell + \text{syntactic}(F')$ for $F' \in \mathcal{P}(\mathcal{T}) \setminus (\{F_{inv}\} \cup \{F_k \mid 1 \leq k \leq 2^\ell - 1\})$ where $\text{syntactic}(F')$ is the syntactic size of F' . Note that size is a generic proof size metric. Then, $\text{CEGVERIF}_{\text{SR}}$ takes $2^{\text{size}(F_{inv})}$ iterations to converge when Abs^{crt} returns the counterexamples $\pi_1, \dots, \pi_{2^\ell - 1}$ in the first $2^\ell - 1$ iterations and Ref returns the proofs $F_1, \dots, F_{2^\ell - 1}$ before returning F_{inv} . \square

B Correspondence of CFGs and CFG Graph Grammar

For an instruction $\text{instr} = v_0 \mapsto (\phi_1 : v_1, \dots, \phi_n : v_n)$, we write $e(\text{instr})$ for $\{(v_0, v_1), \dots, (v_0, v_n)\}$, and $\text{trans}(\text{instr})$ for the map T such that $T((v_0, v_i)) = \phi_i$ for each $i \in \{1, \dots, n\}$. For a term t , we write $e(t)$ for $\bigcup_{\text{instr} \in t} e(\text{instr})$, and $\text{trans}(t)$ for $\bigcup_{\text{instr} \in t} \text{trans}(\text{instr})$.

We define $\text{loops}(t)$ inductively as follows: $\text{loops}(\{\text{instr}\}) = \emptyset$, $\text{loops}(t_1 \cup t_2) = \text{loops}(t_1) \cup \text{loops}(t_2)$, and $\text{loops}(\text{loop } (\mathbf{h}) t) = \text{loops}(t) \cup \{(G, \{\mathbf{h}\})\}$ where $G = (sc(t), e(t) \cap sc(t) \times sc(t))$. Then, a CFG (G, T, v_{ini}, v_{err}) is said to *correspond* to a term t when $v(G) = v(t)$, $e(G) = e(t)$, $T = \text{trans}(t)$, and $\text{loops}(G) = \text{loops}(t)$.

C Formal Definition of cex^{syn}

For a term t and $k > 0$, we define t^k be t with each loop annotation added the unfolding counter k . More formally, we inductively define counter annotation as follows: $\text{instr}^k = \text{instr}$, $(t_1 \cup t_2)^k = t_1^k \cup t_2^k$, and $(\text{loop } (\mathbf{h}) t)^k = \text{loop}^k(\mathbf{h})(t^k)$.

We let the rewriting relation respect the counter annotation, so that each loop is unfolded the number of times it is annotated:

$$\begin{aligned} C[\text{loop}^k(\mathbf{h})\ t] &\dashrightarrow_{ct} C[t\langle v(t)\setminus\{\mathbf{h}\}, sc(t)\rangle \cup \text{loop}^{k-1}(\mathbf{h})\ t\langle\{\mathbf{h}\}, \{\mathbf{h}\}\rangle] \\ C[\text{loop}^0(\mathbf{h})\ t] &\dashrightarrow_{ct} C[t\setminus\{\mathbf{h}\}] \end{aligned}$$

Note that in the first rule, the counter annotations in t are carried over to $t\langle v(t)\setminus\{\mathbf{h}\}, sc(t)\rangle$, and so the nested copies of loops created in the unfolding will also be unfolded the number of times as specified by the annotations in t .

Then, $\text{cex}^{\text{sym}}(P)$ is defined to be the set of CFGs obtained by applying the above reduction to convergence from a term P^k for some $k > 0$. That is, $\text{cex}^{\text{sym}}(P) = \{\pi \mid P^k \dashrightarrow_{ct}^* \pi \wedge k > 0 \wedge \pi \text{ contains no } \text{loop}\}$.

Example. Consider the program P below (left).

<pre> L1: while (b1) { if (b2) { L2: while (b3) { s1; } } else { L3: while (b4) { s2; } } } </pre>	<pre> if (b1) { if (b2) { if (b3) { s1; } if (b3) { s1; } } else { if (b4) { s2; } if (b4) { s2; } } } if (b1) { if (b2) { if (b3) { s1; } if (b3) { s1; } } else { if (b4) { s2; } if (b4) { s2; } } } </pre>
P	π

P has three non-root loops: L_1 , L_2 , and L_3 . The loops L_2 and L_3 are nested inside L_1 . Assuming that the CFG-representation of P preserves the loop structure, the counterexample π obtained by unfolding the loops twice (i.e., $P^2 \dashrightarrow_{ct}^* \pi$) is shown above (right). Note that the outer loop L_1 is unfolded twice, and the inner loops L_2 and L_3 are each unfolded twice per each unfolding of L_1 (and so each four times total).

D Deciding Provability of Acyclic CFGs

We show an algorithm for deciding the provability of acyclic CFGs that is based on a reduction to recursion-free Horn-clause constraint solving [21, 19].

First, we define effective interpolation. Given \mathcal{T} -formulas ϕ and ψ , we say that a \mathcal{T} -formula θ is an *interpolant* of ϕ and ψ if $fv(\theta) \subseteq fv(\phi) \cap fv(\psi)$, $\models_{\mathcal{T}} \phi \Rightarrow \theta$, and $\models_{\mathcal{T}} \theta \Rightarrow \psi$. We say that \mathcal{T} is *effectively interpolating* if there exists an algorithm \mathcal{A} such that, given \mathcal{T} -formulas ϕ and ψ , $\mathcal{A}(\phi, \psi)$ returns an interpolant of ϕ and ψ if one exists, and otherwise returns **No**. By abuse of terminology, for

“quantifier-free theories”, we restrict ϕ, ψ, θ to be quantifier-free formulas of \mathcal{T} in the definition above. QFLRA is known to be effectively interpolating.⁸

Recursion-free Horn-clause Constraints. A *predicate variable application* is of the form $Q(\mathbf{x})$ where Q is a *predicate variable* of arity $|\mathbf{x}|$. A *Horn clause hc* is of the form $\phi \wedge \beta_1 \wedge \dots \wedge \beta_n \Rightarrow \alpha$ where ϕ is a \mathcal{T} -formula, each β_i is a predicate variable application, and α is a predicate variable application or \perp . A *Horn-clause constraint set* (HCCS) \mathcal{H} is a finite set of Horn clauses. We write $pv(\mathcal{H})$ for the predicate variables in \mathcal{H} . We say that \mathcal{H} is *recursion-free* if $\{(Q, Q') \mid \psi \wedge \dots \wedge Q(\mathbf{x}) \dots \Rightarrow Q'(\mathbf{y}) \in \mathcal{H}\}$ is acyclic.

We say that $\rho : pv(\mathcal{H}) \rightarrow \mathcal{T}$ is a *solution* of \mathcal{H} , written $\rho \vdash \mathcal{H}$, if $\models_{\mathcal{T}} \rho(hc)$ for each $hc \in \mathcal{H}$, where $\rho(hc)$ is hc with each predicate variable application $Q(\mathbf{x})$ replaced by $\rho(Q)(\mathbf{x})$. Previous works [21, 19] have shown algorithms for solving recursion-free HCCSs for effectively interpolating \mathcal{T} .⁹

Reduction. We reduce provability of an acyclic CFG to solving a recursion-free HCCS. Let $\gamma = (G, T, v_{ini}, v_{err})$. For each $v \in v(G)$, let Q_v be a distinct predicate variable of arity $|\mathbf{x}|$. Let \mathcal{H}_γ be the set of Horn clauses $\{Q_v(\mathbf{x}) \wedge T((v, v'))(\mathbf{x}, \mathbf{x}') \Rightarrow Q_{v'}(\mathbf{x}') \mid (v, v') \in e(G)\} \cup \{\top \Rightarrow Q_{v_{ini}}(\mathbf{x}), Q_{v_{err}}(\mathbf{x}) \Rightarrow \perp\}$. Because G is acyclic, \mathcal{H}_γ is recursion-free. Also, by construction, we have $\exists \rho : pv(\mathcal{H}_\gamma) \rightarrow F \cup \{\perp, \top\}. \rho \vdash \mathcal{H}_\gamma$ if and only if $F \vdash_{cfg} \gamma$.

E Reducing CFG to one-loop CFG

Let $P = (G, T, v_{ini}, v_{err})$. We give a polynomial time transformation of P into a CFG P^\dagger such that (1) P^\dagger is safe if and only if P is safe, and (2) $cex^{syn}(P^\dagger) = cex(P)$.

For each $v \in v(G)$, let ϕ_v be a \mathcal{T} -predicate of some fixed arity (i.e., $|\mathbf{pc}|$) satisfying $\models_{\mathcal{T}} \phi_v(\mathbf{pc}) \Rightarrow \neg \phi_{v'}(\mathbf{pc})$ for each $v, v' \in v(G)$ such that $v \neq v'$. The transformation works for any theories that have such predicates. For example, for QFLRA, it suffices to let $\phi_v = \lambda pc. pc = num(v)$ where $num(\cdot)$ assigns a distinct rational constant to each $v \in V$.

Let $\mathbf{y} = \mathbf{pc}, \mathbf{x}$ and $\mathbf{y}' = \mathbf{pc}', \mathbf{x}'$ where \mathbf{y}, \mathbf{y}' are distinct variables such that $|\mathbf{y}'| = |\mathbf{y}|$. We define G^\dagger and T^\dagger as follows.

$$\begin{aligned} v(G^\dagger) &= \{v_{ini}^\dagger, v_{err}^\dagger, v_{loop}^\dagger\} \\ e(G^\dagger) &= \{(v_{ini}^\dagger, v_{loop}^\dagger), (v_{loop}^\dagger, v_{loop}^\dagger), (v_{loop}^\dagger, v_{err}^\dagger)\} \\ T^\dagger((v_{ini}^\dagger, v_{loop}^\dagger)) &= \lambda \mathbf{y}, \mathbf{y}'. \phi_{v_{ini}}(\mathbf{pc}') \\ T^\dagger((v_{loop}^\dagger, v_{loop}^\dagger)) &= \lambda \mathbf{y}, \mathbf{y}'. \bigvee_{e \in e(G)} \phi_{sc(e)}(\mathbf{pc}) \wedge T(e)(\mathbf{x}, \mathbf{x}') \wedge \phi_{tg(e)}(\mathbf{pc}') \\ T^\dagger((v_{loop}^\dagger, v_{err}^\dagger)) &= \lambda \mathbf{y}, \mathbf{y}'. \phi_{v_{err}}(\mathbf{pc}) \end{aligned}$$

⁸ This differs from the notion of interpolants from mathematical logic which requires the non-logical symbols of θ to appear in ϕ and ψ . But, this definition is often used in the verification literature and is sufficient for the purpose.

⁹ We remark that the converse also holds, that is, any algorithm for solving recursion-free HCCS gives an interpolation algorithm for the corresponding theory. This is because θ is an interpolant of ϕ and ψ if and only if $\{Q \mapsto \lambda \mathbf{x}. \theta\}$ is a solution of $\{\phi \Rightarrow Q(\mathbf{x}), Q(\mathbf{x}) \wedge \neg \psi \Rightarrow \perp\}$ where $\{\mathbf{x}\} = fv(\phi) \cap fv(\psi)$.

It is easy to see that $P^\dagger = (G^\dagger, T^\dagger, v_{ini}^\dagger, v_{err}^\dagger)$ satisfies the conditions (1) and (2) given above. We note that (2) follows because P^\dagger has only one non-root loop. Also, it can be seen that the syntactic size of the smallest proof of P^\dagger is polynomial in the size of G , the syntactic sizes of ϕ_v 's, and the syntactic size of the smallest proof of P . This follows because $\sigma \vdash_{cfg} P$ implies $\sigma^\dagger \vdash_{cfg} P^\dagger$ where $\sigma^\dagger(v_{ini}^\dagger) = \top$, $\sigma^\dagger(v_{err}^\dagger) = \perp$, and $\sigma^\dagger(v_{loop}^\dagger) = \lambda \mathbf{y}. \bigvee_{v \in v(G)} \phi_v(\mathbf{pc}) \wedge \sigma(v)(\mathbf{x})$.