

SMT-Based Model Checking for Recursive Programs ^{*}

Anvesh Komuravelli, Arie Gurfinkel, and Sagar Chaki

Carnegie Mellon University, Pittsburgh, PA, USA

Abstract. We present an SMT-based symbolic model checking algorithm for safety verification of recursive programs. The algorithm is modular and analyzes procedures individually. Unlike other SMT-based approaches, it maintains both *over-* and *under-approximations* of procedure summaries. Under-approximations are used to analyze procedure calls without inlining. Over-approximations are used to block infeasible counterexamples and detect convergence to a proof. We show that for programs and properties over a decidable theory, the algorithm is guaranteed to find a counterexample, if one exists. However, efficiency depends on an oracle for quantifier elimination (QE). For Boolean Programs, the algorithm is a polynomial decision procedure, matching the worst-case bounds of the best BDD-based algorithms. For Linear Arithmetic (integers and rationals), we give an efficient instantiation of the algorithm by applying QE *lazily*. We use existing interpolation techniques to over-approximate QE and introduce *Model Based Projection* to under-approximate QE. Empirical evaluation on SV-COMP benchmarks shows that our algorithm improves significantly on the state-of-the-art.

1 Introduction

We are interested in the problem of *safety* of recursive programs, *i.e.*, deciding whether an assertion always holds. The first step in Software Model Checking is to approximate the input program by a program model where the program operations are terms in a first-order theory \mathcal{D} . Many program models exist today, e.g., *Boolean Programs* [6] of SLAM [5], GOTO programs of CBMC [14], BOOGIEPL of BOOGIE [7], and, indirectly, internal representations of many tools such as UFO [1], HSF [21], etc. Given a safety property and a program model over \mathcal{D} , it is possible to analyze bounded executions using an oracle for *Satisfiability Modulo Theories* (SMT) for \mathcal{D} . However, in the presence of unbounded recursion, safety is undecidable in general. Throughout this paper, we assume that procedures cannot be passed as parameters.

There exist several program models where safety is efficiently decidable¹, e.g., Boolean Programs with unbounded recursion and the unbounded use of

^{*} This paper is originally published by Springer-Verlag as part of the proceedings of CAV 2014. The final publication is available at link.springer.com.

¹ This is no longer true when we allow procedures as parameters [12].


stack [35,6]. The general observation behind these algorithms is that one can *summarize* the input-output behavior of a procedure. A summary of a procedure is an input-output relation describing what is currently known about its behavior. Thus, a summary can be used to analyze a procedure call without inlining or analyzing the body of the callee [11,36]. For a Boolean Program, the number of states is finite and hence, a summary can only be updated finitely many times. This observation led to a number of efficient algorithms that are polynomial in the number of states, e.g., the RHS framework [35], recursive state machines [4], and symbolic BDD-based algorithms of BEBOP [6] and MOPED [19]. When safety is undecidable (e.g., when \mathcal{D} is Linear Rational Arithmetic (LRA) or Linear Integer Arithmetic (LIA)), several existing software model checkers work by iteratively obtaining Boolean Program abstractions using Predicate Abstraction [13,5]. In this paper, we are interested in an alternative algorithm that works directly on the original program model without an explicit step of Boolean abstraction. Despite the undecidability, we are interested in an algorithm that is guaranteed to find a counterexample to safety, if one exists.

Several algorithms have been recently proposed for verifying recursive programs without predicate abstraction. Notable examples are WHALE [2], HSF [21], GPDR [27], Ultimate Automizer [24,25] and Duality [32]. With the exception of GPDR, these algorithms are based on a combination of Bounded Model Checking (BMC) [8] and Craig Interpolation [16]. First, they use an SMT-solver to check for a bounded counterexample, where the bound is on the depth of the call stack (*i.e.*, the number of nested procedure calls). Second, they use (tree) interpolation to over-approximate procedure summaries. This is repeated with increasing values of the bound until a counterexample is found or the approximate summaries are inductive. The reduction to BMC ensures that the algorithms are guaranteed to find a counterexample. However, the size of the SMT instance grows exponentially with the bound on the call-stack (*i.e.*, linear in the size of the call tree). Therefore, for Boolean Programs, these algorithms are at least worst-case exponential in the number of states.

On the other hand, GPDR follows the approach of IC3 [9] by solving BMC incrementally without unrolling the call-graph. Interpolation is used to over-approximate summaries and caching is used to indirectly under-approximate them. For some configurations, GPDR is worst-case polynomial for Boolean Programs. However, even for LRA, GPDR might fail to find a counterexample.²

In this paper, we introduce RECMC, the first SMT-based algorithm for model checking safety of recursive programs that is worst-case polynomial (in the number of states) for Boolean Programs while being a co-semidecision procedure for programs over decidable theories (see Section 4). Our main insight is to maintain not only over-approximations of procedure summaries (which we call *summary facts*), but also their under-approximations (which we call *reachability facts*). While summary facts are used to block spurious counterexamples, reachability facts are used to analyze a procedure call without inlining or analyzing the body of the callee. Our use of reachability facts is similar to that of *summary edges*

² See appendix for an example.

Fig. 1: Flow of the algorithm RECMC to check if $M \models \varphi_{safe}$.

of the RHS [35] algorithm. This explains our complexity result for Boolean Programs. However, our summary facts make an important difference. While the use of summary facts is an interesting heuristic for Boolean Programs that does not improve the worst-case complexity, it is crucial for richer theories.


Almost every step of RECMC results in existential quantification of variables. RECMC tries to eliminate these variables, as otherwise, they would accumulate and the size of an inferred reachability fact, for example, grows exponentially in the bound on the call-stack. But, a naïve use of quantifier elimination (QE) is expensive. Instead, we develop an alternative approach that under-approximates QE. However, obtaining arbitrary under-approximations can lead to divergence of the algorithm. We introduce the concept of *Model Based Projection* (MBP), for *covering* $\exists \bar{x} \cdot \varphi(\bar{x}, \bar{y})$ by *finitely-many* quantifier-free under-approximations obtained using models of $\varphi(\bar{x}, \bar{y})$. We developed efficient MBPs (see Section 5) for Linear Arithmetic based on the QE methods by Loos-Weispfenning [30] for LRA and Cooper [15] for LIA. We use MBP to under-approximate reachability facts in RECMC. In the best case, only a partial under-approximation is needed and a complete quantifier elimination can be avoided.

We have implemented RECMC as part of our tool SPACER using the framework of Z3 [17] and evaluated it on 799 benchmarks from SV-COMP [37]. SPACER significantly outperforms the implementation of GPDR in Z3 (see Section 6).

In summary, our contributions are: (a) an efficient SMT-based algorithm for model checking recursive programs, that analyzes procedures individually using under- and over-approximations of procedure summaries, (b) MBP functions for under-approximating quantifier elimination for LRA and LIA, (c) a new, complete algorithm for Boolean Programs, with complexity polynomial in the number of states, similar to the best known method [6], and (d) an implementation and an empirical evaluation of the approach.

2 Overview

In this section, we give an overview of RECMC and illustrate it on an example. Let \mathcal{A} be a recursive program. For simplicity of presentation, assume no loops, no global variables and that arguments are passed by reference. Let $P(\bar{v}) \in \mathcal{A}$ be a procedure with parameters \bar{v} and let \bar{v}_0 be fresh variables not appearing in P with $|\bar{v}| = |\bar{v}_0|$. A safety property for P is an assertion $\varphi(\bar{v}_0, \bar{v})$. We say that P satisfies φ , denoted $P(\bar{v}) \models \varphi(\bar{v}_0, \bar{v})$, iff the Hoare-triple $\{\bar{v} = \bar{v}_0\} P(\bar{v}) \{\varphi(\bar{v}_0, \bar{v})\}$ is valid.

Fig. 2: Flow of the algorithm BND SAFETY to check $P \models_b \varphi$.

Note that every Hoare-triple corresponds to a safety property in this sense, as shown by Clarke [11], using a *Rule of Adaptation*. Given a safety property φ and a natural number $n \geq 0$, the problem of *bounded safety* is to determine whether all executions of P using a call-stack bounded by n satisfy φ . We use $P(\bar{v}) \models_n \varphi(\bar{v}_0, \bar{v})$ to denote bounded safety.

The key steps of RECMC are shown in Fig. 1. RECMC decides safety for the main procedure M of \mathcal{A} . RECMC maintains two *assertion maps* ρ and σ . The *reachability* map ρ maps each procedure $P(\bar{v}) \in \mathcal{A}$ to a set of assertions over $\bar{v}_0 \cup \bar{v}$ that under-approximate its behavior. Similarly, the *summary* map σ maps a procedure P to a set of assertions that over-approximate its behavior. Given P , the maps are partitioned according to the bound on the call-stack. That is, if $\delta(\bar{v}_0, \bar{v}) \in \rho(P, n)$ for $n \geq 0$, then for every model m of δ , there is an execution of P that begins in $m(\bar{v}_0)$ and ends in $m(\bar{v})$, using a call-stack bounded by n . Similarly, if $\delta(\bar{v}_0, \bar{v}) \in \sigma(P, n)$, then $P(\bar{v}) \models_n \delta(\bar{v}_0, \bar{v})$.

RECMC alternates between two steps: **(A)** deciding bounded safety (that also updates ρ and σ maps) and **(B)** checking whether the current proof of bounded safety is inductive (*i.e.*, independent of the bound). It terminates when a counterexample or a proof is found.

Bounded safety, $P \models_b \varphi$, is decided using BND SAFETY shown in Fig. 2. Step 1 checks whether φ is falsified by current reachability facts in ρ of the callees of P . If so, it infers a new reachability fact for P at bound b witnessing the falsification of φ . Step 2 checks whether φ is satisfied using current summary facts in σ of the callees. If so, it infers a new summary fact for P at bound b witnessing the satisfaction of φ . If the prior two steps fail, there is a potential counterexample π in P with a call to some procedure R such that the reachability facts of R are too strong to witness π , but the summary facts of R are too weak to block it. Step 3 updates ρ and σ by creating (and recursively deciding) a new bounded safety problem for R at bound $b - 1$.

We conclude this section with an illustration of RECMC on the program in Fig. 3 (adapted from [11]). The program has 3 procedures: the main procedure M , and procedures T and D . M calls T and D . T modifies its argument t and calls itself recursively. D decrements its argument d . Let the property be $\varphi = m_0 \geq 2m + 4$.

The first iteration of RECMC is trivial. The bound $n = 0$ and since M has no call-free executions it vacuously satisfies any bounded safety property. Fig. 4 shows the four iterations of BND SAFETY for the second iteration of RECMC where $n = 1$. For this bound, the maps ρ and σ are initially empty. The first




Fig. 3: A recursive program with 3 procedures.

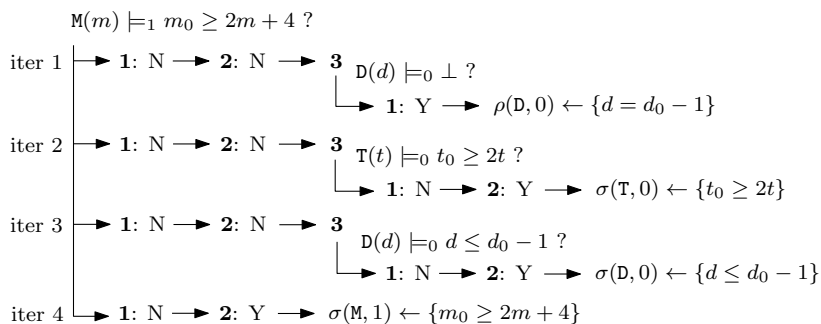


Fig. 4: A run of BNDSAFETY on program in Fig. 3 and a bound 1 on the stack depth. Numbers in bold refer to the steps in Fig. 2.

iteration of BNDSAFETY finds a potential counterexample path in M and the approximation for D is updated with a new reachability fact: $d = d_0 - 1$. In the second iteration, the approximation for T is updated. Note that the two calls to D are “jumped over” using the reachability fact for D computed in the first iteration. The new summary fact for T is: $t_0 \geq 2t$. In the third iteration, the approximation for D is updated again, now with a summary fact $d \leq d_0 - 1$. Finally, the summary facts for T and D at bound 0 are sufficient to establish bounded safety at $n = 1$. At this point, the summary map σ is:

$$\sigma(M, 1) = \{m_0 \geq 2m + 4\} \quad \sigma(T, 0) = \{t_0 \geq 2t\} \quad \sigma(D, 0) = \{d \leq d_0 - 1\}$$

Ignoring the bounds, σ is inductive. For example, we can prove that the body of T satisfies $t_0 \geq 2t$, assuming that the calls do. Thus, step **B** of RECMC succeeds and the algorithm terminates declaring the program SAFE. In the rest of the paper, we show how to automate RECMC using an SMT-oracle.

3 Preliminaries

Consider a first-order language with equality and let \mathcal{S} be its signature, *i.e.*, the set of non-logical function and predicate symbols (including equality). An \mathcal{S} -structure I consists of a domain of interpretation, denoted $|I|$, and assigns elements of $|I|$ to variables, and functions and predicates on $|I|$ to the symbols of \mathcal{S} . Let φ be a formula. We assume the usual definition of satisfaction of φ by I , denoted $I \models \varphi$. I is called a *model* of φ iff $I \models \varphi$ and this can be extended to a set of formulas. A first-order \mathcal{S} -theory Th is a set of deductively closed \mathcal{S} -sentences. I satisfies φ modulo Th , denoted $I \models_{Th} \varphi$, iff $I \models Th \cup \{\varphi\}$. φ is *valid* modulo Th , denoted $\models_{Th} \varphi$, iff every model of Th is also a model of φ .

Let I be an \mathcal{S} -structure and \bar{w} be a list of fresh function/predicate symbols not in \mathcal{S} . A $(\mathcal{S} \cup \bar{w})$ -structure J is called an *expansion* of I to \bar{w} iff $|J| = |I|$ and J agrees with I on the assignments to all variables and the symbols of \mathcal{S} . We use the notation $I\{\bar{w} \mapsto \bar{u}\}$ to denote the expansion of I to \bar{w} that assigns the function/predicate u_i to the symbol w_i . For an \mathcal{S} -sentence φ , we write $I(\varphi)$ to denote the truth value of φ under I . For a formula $\varphi(\bar{x})$ with free variables \bar{x} , we overload the notation $I(\varphi)$ to mean $\{\bar{a} \in |I|^{|\bar{x}|} \mid I\{\bar{x} \mapsto \bar{a}\} \models \varphi\}$. For simplicity of presentation, we sometimes identify the truth value *true* with $|I|$ and *false* with \emptyset .

We assume that programs do not have internal procedures and that procedures cannot be passed as parameters. Furthermore, without loss of generality, we assume that programs do not have loops or global variables. In the following, we define programs using a logical representation, as opposed to giving a concrete syntax. A *program* \mathcal{A} is a finite list of procedures with a designated *main* procedure M where the program begins. A *procedure* P is a tuple $\langle \bar{l}_P, \bar{o}_P, \Sigma_P, \bar{\ell}_P, \beta_P \rangle$, where (a) \bar{l}_P is the finite list of variables denoting the input values of the parameters, (b) \bar{o}_P is the finite list of variables denoting the output values of the parameters, (c) Σ_P is a fresh predicate symbol of arity $|\bar{l}_P| + |\bar{o}_P|$, (d) $\bar{\ell}_P$ is the finite list of local variables, and (e) β_P is a quantifier-free sentence over the signature $(\mathcal{S} \cup \{\Sigma_Q \mid Q \in \mathcal{A}\} \cup \bar{l}_P \cup \bar{o}_P \cup \bar{\ell}_P)$ in which a predicate symbol Σ_Q appears only positively. We use \bar{v}_P to denote $\bar{l}_P \cup \bar{o}_P$.

Intuitively, for a procedure P , Σ_P is used to denote its semantics and β_P encodes its body using the predicate symbol Σ_Q for a call to the procedure Q . We require that a predicate symbol Σ_Q appears only positively in β_P to ensure a fixed-point characterization of the semantics as shown later on. For example, for the signature $\mathcal{S} = \langle 0, Succ, -, +, \leq, >, = \rangle$, the program in Fig. 3 is represented as $\langle M, T, D \rangle$ with $M = \langle m_0, m, \Sigma_M, \langle \ell_0, \ell_1 \rangle, \beta_M \rangle$, $T = \langle t_0, t, \Sigma_T, \langle \ell_0, \ell_1 \rangle, \beta_T \rangle$ and $D = \langle d_0, d, \Sigma_D, \emptyset, \beta_D \rangle$, where

$$\begin{aligned} \beta_M &= \Sigma_T(m_0, \ell_0) \wedge \Sigma_D(\ell_0, \ell_1) \wedge \Sigma_D(\ell_1, m) & \beta_D &= (d = d_0 - 1) \\ \beta_T &= (t_0 \leq 0 \wedge t_0 = t) \vee (t_0 > 0 \wedge \ell_0 = t_0 - 2 \wedge \Sigma_T(\ell_0, \ell_1) \wedge t = \ell_1 + 1) \end{aligned} \quad (1)$$

Here, we abbreviate $Succ^i(0)$ by i and (m_0, t_0, d_0) and (m, t, d) denote the input and the output values of the parameters of the original program, respectively. For a procedure P , let $Paths(P)$ denote the set of all prime-implicants of β_P . Intuitively, each element of $Paths(P)$ encodes a path in the procedure.

Let $\mathcal{A} = \langle P_0, \dots, P_n \rangle$ be a program and I be an \mathcal{S} -structure. Let \bar{X} be a list of length n such that each X_i is either (i) a truth value if $|\bar{v}_{P_i}| = 0$, or (ii) a subset of $|I|^{|\bar{v}_{P_i}|}$ if $|\bar{v}_{P_i}| \geq 1$. Let $J(I, \bar{X})$ denote the expansion $I\{\Sigma_{P_0} \mapsto X_0\} \dots \{\Sigma_{P_n} \mapsto X_n\}$. The *semantics* of a procedure P_i given I , denoted $\llbracket P_i \rrbracket_I$, characterizes all the terminating executions of P_i and is defined as follows. $\langle \llbracket P_0 \rrbracket_I, \dots, \llbracket P_n \rrbracket_I \rangle$ is the (pointwise) least \bar{X} such that for all $Q \in \mathcal{A}$, $J(I, \bar{X}) \models \forall \bar{v}_Q \cup \bar{\ell}_Q \cdot (\beta_Q \Rightarrow \Sigma_Q(\bar{v}_Q))$. This has a well-known least fixed-point characterization [11].

For a bound $b \geq 0$ on the call-stack, the *bounded semantics* of a procedure P_i given I , denoted $\llbracket P_i \rrbracket_I^b$, characterizes all the executions using a stack of depth

bounded by b and is defined by induction on b :

$$\llbracket P_i \rrbracket_I^0 = J(I, \langle \emptyset, \dots, \emptyset \rangle) (\exists \bar{\ell}_{P_i} \cdot \beta_{P_i}), \quad \llbracket P_i \rrbracket_I^b = J(I, \langle \llbracket P_0 \rrbracket_I^{b-1}, \dots, \llbracket P_n \rrbracket_I^{b-1} \rangle) (\exists \bar{\ell}_{P_i} \cdot \beta_{P_i})$$

An *environment* is a function that maps a predicate symbol Σ_P to a formula over \bar{v}_P . Given a formula τ and an environment E , we abuse the notation $\llbracket \cdot \rrbracket$ and write $\llbracket \tau \rrbracket_E$ for the formula obtained by instantiating every predicate symbol Σ_P by $E(\Sigma_P)$ in τ .

Let Th be an \mathcal{S} -theory. A *safety property* for a procedure $P \in \mathcal{A}$ is a formula over \bar{v}_P . P satisfies a safety property φ w.r.t Th , denoted $P \models_{Th} \varphi$, iff for all models I of Th , $\llbracket P \rrbracket_I \subseteq I(\varphi)$. A *safety property* ψ of the program \mathcal{A} is a safety property of its main procedure. A *safety proof* for $\psi(\bar{v}_M)$ is an environment Π that is both safe and inductive:

$$\models_{Th} \llbracket \forall \bar{x} \cdot \Sigma_M(\bar{x}) \Rightarrow \psi(\bar{x}) \rrbracket_\Pi, \quad \forall P \in \mathcal{A} \models_{Th} \llbracket \forall \bar{v}_P \cup \bar{\ell}_P \cdot (\beta_P \Rightarrow \Sigma_P(\bar{v}_P)) \rrbracket_\Pi$$

Given a formula $\varphi(\bar{v}_P)$ and $b \geq 0$, a procedure P satisfies *bounded safety* w.r.t Th , denoted $P \models_{b,Th} \varphi$, iff for all models I of Th , $\llbracket P \rrbracket_I^b \subseteq I(\varphi)$. In this case, we also call φ a *summary fact* for $\langle P, b \rangle$. We call φ a *reachability fact* for $\langle P, b \rangle$ iff $I(\varphi) \subseteq \llbracket P \rrbracket_I^b$, for all models I of Th . Intuitively, *summary facts* and *reachability facts* for $\langle P, b \rangle$, respectively, over- and under-approximate $\llbracket P \rrbracket_I^b$ for every model I of Th .

A *bounded assertion map* maps a procedure P and a natural number $b \geq 0$ to a set of formulas over \bar{v}_P . Given a bounded assertion map m and $b \geq 0$, we define two special environments U_m^b and O_m^b as follows.

$$U_m^b : \Sigma_P \mapsto \bigvee \{ \delta \in m(P, b') \mid b' \leq b \} \quad O_m^b : \Sigma_P \mapsto \bigwedge \{ \delta \in m(P, b') \mid b' \geq b \}$$

We use U_m^b and O_m^b to under- and over-approximate the bounded semantics. For convenience, let U_m^{-1} and O_m^{-1} be environments that map every symbol to \perp .

4 Model Checking Recursive Programs

In this section, we present our algorithm $\text{RECMC}(\mathcal{A}, \varphi_{safe})$ that determines whether a program \mathcal{A} satisfies a safety property φ_{safe} . Let \mathcal{S} be the signature of the first-order language under consideration and assume a fixed \mathcal{S} -theory Th . To avoid clutter, we drop the subscript Th from the notation \models_{Th} and $\models_{b,Th}$. We also establish the soundness and complexity of RECMC. An efficient instantiation of RECMC to Linear Arithmetic is presented in Section 5.

Main Loop. RECMC maintains two *bounded assertion maps* ρ and σ for reachability and summary facts, respectively. For brevity, for a first-order formula τ , we write $\llbracket \tau \rrbracket_\rho^b$ and $\llbracket \tau \rrbracket_\sigma^b$ to denote $\llbracket \tau \rrbracket_{U_\rho^b}$ and $\llbracket \tau \rrbracket_{O_\sigma^b}$, respectively, where the environments U_ρ^b and O_σ^b are as defined in Section 3. Intuitively, $\llbracket \tau \rrbracket_\rho^b$ and $\llbracket \tau \rrbracket_\sigma^b$, respectively, under- and over-approximate τ using ρ and σ .

```

RECMC( $\mathcal{A}, \varphi_{safe}$ )
1  $n \leftarrow 0; \rho \leftarrow \emptyset; \sigma \leftarrow \emptyset$ 
2 while true do
3    $res, \rho, \sigma \leftarrow \text{BND\_SAFETY}(\mathcal{A}, \varphi_{safe}, n, \rho, \sigma)$ 
4   if  $res$  is UNSAFE then
5      $\perp$  return UNSAFE,  $\rho$ 
6   else
7      $ind, \sigma \leftarrow \text{CHECK\_INDUCTIVE}(\mathcal{A}, \sigma, n)$ 
8     if  $ind$  then
9        $\perp$  return SAFE,  $\sigma$ 
9      $n \leftarrow n + 1$ 

CHECKINDUCTIVE( $\mathcal{A}, \sigma, n$ )
10  $ind \leftarrow true$ 
11 foreach  $P \in \mathcal{A}$  do
12   foreach  $\delta \in \sigma(P, n)$  do
13     if  $\models \llbracket \beta_P \rrbracket_{\sigma}^n \Rightarrow \delta$  then
14        $\perp$   $\sigma \leftarrow \sigma \cup (\langle P, n + 1 \rangle \mapsto \delta)$ 
15     else
16        $\perp$   $ind \leftarrow false$ 
16 return ( $ind, \sigma$ )

```

Fig. 5: Pseudo-code of RECMC.

The pseudo-code of the main loop of RECMC (corresponding to the flow diagram in Fig. 1) is shown in Fig. 5. RECMC follows an *iterative deepening* strategy. In each iteration, BND SAFETY (described below) checks whether all executions of \mathcal{A} satisfy φ_{safe} for a bound $n \geq 0$ on the call-stack, *i.e.*, if $M \models_n \varphi_{safe}$. BND SAFETY also updates the maps ρ and σ . Whenever BND SAFETY returns *UNSAFE*, the reachability facts in ρ are sufficient to construct a counterexample and the loop terminates. Whenever BND SAFETY returns *SAFE*, the summary facts in σ are sufficient to prove the absence of a counterexample for the current bound n on the call-stack. In this case, if σ is also inductive, as determined by CHECKINDUCTIVE, O_{σ}^n is a safety proof and the loop terminates. Otherwise, the bound on the call-stack is incremented and a new iteration of the loop begins. Note that, as a side-effect of CHECKINDUCTIVE, some summary facts are propagated to the bound $n + 1$. This is similar to *push generalization* in IC3 [9].

Bounded safety. We describe the routine $\text{BND_SAFETY}(\mathcal{A}, \varphi_{safe}, n, \rho_{Init}, \sigma_{Init})$ as an abstract transition system [33] defined by the inference rules shown in Fig. 6. Here, n is the current bound on the call-stack and ρ_{Init} and σ_{Init} are the maps of reachability and summary facts input to the routine. A state of BND SAFETY is a triple $\mathcal{Q} \parallel \rho \parallel \sigma$, where ρ and σ are the current maps and \mathcal{Q} is a set of triples $\langle P, \varphi, b \rangle$ for a procedure P , a formula φ over \bar{v}_P , and a number $b \geq 0$. A triple $\langle P, \varphi, b \rangle \in \mathcal{Q}$ is called a *bounded reachability query* and asks whether $P \not\models_b \neg\varphi$, *i.e.*, whether there is an execution in P using a call-stack bounded by b where the values of \bar{v}_P satisfy φ .

BND SAFETY starts with a single query $\langle M, \neg\varphi_{safe}, n \rangle$ and initializes the maps of reachability and summary facts (rule INIT). It checks whether $M \models_n \varphi_{safe}$ by inferring new summary and reachability facts to answer existing queries (rules SUM and REACH) and generating new queries (rule QUERY). When there are no queries left to answer, *i.e.*, \mathcal{Q} is empty, it terminates with a result of either *UNSAFE* or *SAFE* (rules UNSAFE and SAFE).

SUM infers a new summary fact when a query $\langle P, \varphi, b \rangle$ can be answered negatively. In this case, there is an over-approximation of the bounded semantics of P at b , obtained using the summary facts of callees at bound $b - 1$, that is unsatisfiable with φ . That is, $\models \llbracket \beta_P \rrbracket_{\sigma}^{b-1} \Rightarrow \neg\varphi$. The inference of the new fact is by interpolation [16] (denoted by ITP in the side-condition of the rule). Thus, the new summary fact ψ is a formula over \bar{v}_P such that $\models (\llbracket \beta_P \rrbracket_{\sigma}^{b-1} \Rightarrow \psi(\bar{v}_P)) \wedge (\psi(\bar{v}_P) \Rightarrow \neg\varphi)$. Note that ψ over-approximates the bounded semantics of P at b . Every

$$\begin{array}{c}
\text{INIT} \frac{}{\{ \langle M, \neg \varphi_{safe}, n \rangle \} \parallel \rho_{Init} \parallel \sigma_{Init}} \\
\\
\text{SUM} \frac{\mathcal{Q} \parallel \rho \parallel \sigma \quad \langle P, \varphi, b \rangle \in \mathcal{Q} \quad \models \llbracket \beta_P \rrbracket_{\sigma}^{b-1} \Rightarrow \neg \varphi}{\mathcal{Q} \setminus \{ \langle P, \eta, c \rangle \mid c \leq b, \models \llbracket \Sigma_P \rrbracket_{\sigma}^c \wedge \psi \Rightarrow \neg \eta \} \parallel \rho \parallel \sigma \cup \{ \langle P, b \rangle \mapsto \psi \}} \\
\text{where } \psi = \text{ITP}(\llbracket \beta_P \rrbracket_{\sigma}^{b-1}, \neg \varphi) \\
\\
\text{REACH} \frac{\mathcal{Q} \parallel \rho \parallel \sigma \quad \langle P, \varphi, b \rangle \in \mathcal{Q} \quad \pi \in \text{Paths}(P) \quad \not\models \llbracket \pi \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi}{\mathcal{Q} \setminus \{ \langle P, \eta, c \rangle \mid c \geq b, \not\models \psi \Rightarrow \neg \eta \} \parallel \rho \cup \{ \langle P, b \rangle \mapsto \psi \} \parallel \sigma} \\
\text{where } \psi = \exists \bar{\ell}_P \cdot \llbracket \pi \rrbracket_{\rho}^{b-1} \\
\\
\text{QUERY} \frac{\mathcal{Q} \parallel \rho \parallel \sigma \quad \langle P, \varphi, b \rangle \in \mathcal{Q} \quad \models \llbracket \beta_P \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi \quad \pi \in \text{Paths}(P) \\ \pi = \pi_u \wedge \Sigma_R(\bar{a}) \wedge \pi_v \quad \models \llbracket \pi_u \rrbracket_{\sigma}^{b-1} \wedge \llbracket \Sigma_R(\bar{a}) \rrbracket_{\rho}^{b-1} \wedge \llbracket \pi_v \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi \\ \not\models \llbracket \pi_u \rrbracket_{\sigma}^{b-1} \wedge \llbracket \Sigma_R(\bar{a}) \rrbracket_{\sigma}^{b-1} \wedge \llbracket \pi_v \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi}{\mathcal{Q} \cup \{ \langle R, \psi, b-1 \rangle \} \parallel \rho \parallel \sigma} \\
\text{where } \begin{cases} \psi = (\exists (\bar{v}_P \cup \bar{\ell}_P) \setminus \bar{a} \cdot \llbracket \pi_u \rrbracket_{\sigma}^{b-1} \wedge \llbracket \pi_v \rrbracket_{\rho}^{b-1} \wedge \varphi) [\bar{a} \leftarrow \bar{v}_R] \\ \text{for all } \langle R, \eta, b-1 \rangle \in \mathcal{Q}, \models \psi \Rightarrow \neg \eta \end{cases} \\
\\
\text{UNSAFE} \frac{\emptyset \parallel \rho \parallel \sigma \quad \not\models \llbracket \Sigma_M \rrbracket_{\rho}^n \Rightarrow \varphi_{safe}}{\text{UNSAFE}} \quad \text{SAFE} \frac{\emptyset \parallel \rho \parallel \sigma \quad \models \llbracket \Sigma_M \rrbracket_{\sigma}^n \Rightarrow \varphi_{safe}}{\text{SAFE}}
\end{array}$$

Fig. 6: Rules defining $\text{BND_SAFETY}(\mathcal{A}, \varphi_{safe}, n, \rho_{Init}, \sigma_{Init})$.

query $\langle P, \eta, c \rangle \in \mathcal{Q}$ such that η is unsatisfiable with the updated environment $O_{\sigma}^c(\Sigma_P)$ is immediately answered and removed.

REACH infers a new reachability fact when a query $\langle P, \varphi, b \rangle$ can be answered positively. In this case, there is an under-approximation of the bounded semantics of P at b , obtained using the reachability facts of callees at bound $b-1$, that is satisfiable with φ . That is, $\not\models \llbracket \beta_P \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi$. In particular, there exists a path π in $\text{Paths}(P)$ such that $\not\models \llbracket \pi \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi$. The new reachability fact ψ is obtained by choosing such a π non-deterministically and existentially quantifying all local variables from $\llbracket \pi \rrbracket_{\rho}^{b-1}$. Note that ψ under-approximates the bounded semantics of P at b . Every query $\langle P, \eta, c \rangle \in \mathcal{Q}$ such that η is satisfiable with the updated environment $U_{\rho}^c(\Sigma_P)$ is immediately answered and removed.

QUERY creates a new query when a query $\langle P, \varphi, b \rangle$ cannot be answered using current ρ and σ . In this case, the current over-approximation of the bounded semantics of P at b is satisfiable with φ while its current under-approximation is unsatisfiable with φ . That is, $\not\models \llbracket \beta_P \rrbracket_{\sigma}^{b-1} \Rightarrow \neg \varphi$ and $\models \llbracket \beta_P \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi$. In particular, there exists a path π in $\text{Paths}(P)$ such that $\not\models \llbracket \pi \rrbracket_{\sigma}^{b-1} \Rightarrow \neg \varphi$ and $\models \llbracket \pi \rrbracket_{\rho}^{b-1} \Rightarrow \neg \varphi$. Intuitively, π is a potential counterexample path that needs to be checked for feasibility. Such a π is chosen non-deterministically. π is guaranteed to have a call $\Sigma_R(\bar{a})$ to a procedure R such that the under-approximation $\llbracket \Sigma_R(\bar{a}) \rrbracket_{\rho}^{b-1}$ is too strong to witness π but the over-approximation $\llbracket \Sigma_R(\bar{a}) \rrbracket_{\sigma}^{b-1}$ is

	π_i	$\llbracket \pi_i \rrbracket_\rho^0$	$\llbracket \pi_i \rrbracket_\sigma^0$
$i = 1$	$\Sigma_T(m_0, \ell_0)$	\perp	\top
$i = 2$	$\Sigma_D(\ell_0, \ell_1)$	$\ell_1 = \ell_0 - 1$	\top
$i = 3$	$\Sigma_D(\ell_1, m)$	$m = \ell_1 - 1$	\top

Fig. 7: Approximations of the only path π of the procedure M in Fig. 3.

too weak to block it. That is, π can be partitioned into a prefix π_u , a call $\Sigma_R(\bar{a})$ to R , and a suffix π_v such that the following hold:

$$\begin{aligned} & \models \llbracket \Sigma_R(\bar{a}) \rrbracket_\rho^{b-1} \Rightarrow ((\llbracket \pi_u \rrbracket_\sigma^{b-1} \wedge \llbracket \pi_v \rrbracket_\rho^{b-1}) \Rightarrow \neg\varphi) \\ & \not\models \llbracket \Sigma_R(\bar{a}) \rrbracket_\sigma^{b-1} \Rightarrow ((\llbracket \pi_u \rrbracket_\sigma^{b-1} \wedge \llbracket \pi_v \rrbracket_\rho^{b-1}) \Rightarrow \neg\varphi) \end{aligned}$$

Note that the prefix π_u and the suffix π_v are over- and under-approximated, respectively. A new query $\langle R, \psi, b-1 \rangle$ is created where ψ is obtained by existentially quantifying all variables from $\llbracket \pi_u \rrbracket_\sigma^{b-1} \wedge \llbracket \pi_v \rrbracket_\rho^{b-1} \wedge \varphi$ except the arguments \bar{a} of the call, and renaming appropriately. If the new query is answered negatively (using SUM), all executions along π where the values of $\bar{v}_P \cup \bar{\ell}_P$ satisfy $\llbracket \pi_v \rrbracket_\rho^{b-1}$ are spurious counterexamples. An additional side-condition requires that ψ “does not overlap” with η for any other query $\langle R, \eta, b-1 \rangle$ in \mathcal{Q} . This is necessary for termination of BNDSAFETY (Theorem 2). In practice, the side-condition is trivially satisfied by always applying the rule to $\langle P, \varphi, b \rangle$ with the smallest b .

For example, consider the program in Fig. 3 represented by (1) and the query $\langle M, \varphi, 1 \rangle$ where $\varphi \equiv m_0 < 2m + 4$. Let $\sigma = \emptyset$, $\rho(D, 0) = \{d = d_0 - 1\}$ and $\rho(T, 0) = \emptyset$. Let $\pi = (\Sigma_T(m_0, \ell_0) \wedge \Sigma_D(\ell_0, \ell_1) \wedge \Sigma_D(\ell_1, m))$ denote the only path in the procedure M . Fig. 7 shows $\llbracket \pi_i \rrbracket_\rho^0$ and $\llbracket \pi_i \rrbracket_\sigma^0$ for each conjunct π_i of π . As the figure shows, $\llbracket \pi \rrbracket_\sigma^0$ is satisfiable with φ , witnessed by the execution $e \equiv \langle m_0 = 3, \ell_0 = 3, \ell_1 = 2, m = 1 \rangle$. Note that this execution also satisfies $\llbracket \pi_2 \wedge \pi_3 \rrbracket_\rho^0$. But, $\llbracket \pi_1 \rrbracket_\rho^0$ is too strong to witness it, where π_1 is the call $\Sigma_T(m_0, \ell_0)$. To create a new query for T , we first existentially quantify all variables other than the arguments m_0 and ℓ_0 from $\pi_2 \wedge \pi_3 \wedge \varphi$, obtaining $m_0 < 2\ell_0$. Renaming the arguments by the parameters of T results in the new query $\langle T, t_0 < 2t, 0 \rangle$. Further iterations of BNDSAFETY would answer this query negatively making the execution e spurious. Note that this would also make all other executions where the values to $\langle m_0, \ell_0, \ell_1, m \rangle$ satisfy $\llbracket \pi_2 \wedge \pi_3 \rrbracket_\rho^0$ spurious.

Soundness and Complexity. Soundness of RECMC follows from that of BNDSAFETY, which can be shown by a case analysis on the inference rules³.

Theorem 1. BNDSAFETY and RECMC are sound.

BNDSAFETY is complete relative to an oracle for satisfiability modulo Th . Even though the number of reachable states of a procedure is unbounded in general, the number of reachability facts inferred by BNDSAFETY is finite. This is because a reachability fact corresponds to a path (see REACH) and given a bound on the call-stack, the number of such facts is bounded. This further bounds the number of queries that can be created.

³ Proofs of all of the theorems are in the Appendix.

Theorem 2. *Given an oracle for Th , $\text{BNDSAFETY}(\mathcal{A}, \varphi, n, \emptyset, \emptyset)$ terminates.*

As a corollary of Theorem 2, RECMC is a co-semidecision procedure for safety, *i.e.*, RECMC is guaranteed to find a counterexample if one exists. In contrast, the closest related algorithm GPDR [27] is not a co-semidecision procedure (see Appendix). Finally, for Boolean Programs RECMC is a complete decision procedure. Unlike the general case, the number of reachable states of a Boolean Program, and hence the number of reachability facts, is finite and independent of the bound on the call-stack. Let $N = |\mathcal{A}|$ and $k = \max\{|\bar{v}_P| \mid P \in \mathcal{A}\}$.

Theorem 3. *Let \mathcal{A} be a Boolean Program. Then $\text{RECMC}(\mathcal{A}, \varphi)$ terminates in $O(N^2 \cdot 2^{2k})$ -many applications of the rules in Fig. 6.*

Note that due to the iterative deepening strategy of RECMC, the complexity is quadratic in the number of procedures (and not linear as in [6]). In contrast, other SMT-based algorithms, such as WHALE [2], are worst-case exponential in the number of states of a Boolean Program.

In summary, RECMC checks safety of a recursive program by inferring the necessary under- and over-approximations of procedure semantics and using them to analyze procedures individually.

5 Model Based Projection

RECMC, as presented in Section 4, can be used as-is, when Th is Linear Arithmetic. But, note that the rules REACH and QUERY introduce existential quantifiers in reachability facts and queries. Unless eliminated, these quantifiers accumulate and the size of the formulas grows exponentially in the bound on the call-stack. Using quantifier elimination (QE) to eliminate the quantifiers is expensive. Instead, we suggest an alternative approach that under-approximates existential quantification with quantifier-free formulas *lazily* and efficiently. We first introduce a model-based under-approximation of QE, which we call a *Model Based Projection* (MBP). Second, we give an efficient (linear in the size of formulas involved) MBP procedure for Linear Rational Arithmetic (LRA). Due to space limitations, MBP for Linear Integer Arithmetic (LIA) is described in the Appendix. Finally, we show a modified version of BNDSAFETY that uses MBP instead of existential quantification and show that it is sound and terminating. **Model Based Projection (MBP).** Let $\lambda(\bar{y})$ be the existentially quantified formula $\exists \bar{x} \cdot \lambda_m(\bar{x}, \bar{y})$ where λ_m is quantifier free. A function Proj_λ from models (modulo Th) of λ_m to quantifier-free formulas over \bar{y} is a *Model Based Projection* (for λ) iff it has a finite image, $\lambda \equiv \bigvee_{M \models \lambda_m} \text{Proj}_\lambda(M)$, and for every model M of λ_m , $M \models \text{Proj}_\lambda(M)$.

In other words, Proj_λ covers the space of all models of $\lambda_m(\bar{x}, \bar{y})$ by a finite set of quantifier-free formulas over \bar{y} . MBP exists for any theory that admits quantifier elimination, because one can first obtain an equivalent quantifier-free formula and map every model to it.

MBP for Linear Rational Arithmetic. We begin with a brief overview of Loos-Weispfenning (LW) method [30] for quantifier elimination in LRA. We borrow our presentation from Nipkow [34] to which we refer the reader for more details. Let $\lambda(\bar{y}) = \exists \bar{x} \cdot \lambda_m(\bar{x}, \bar{y})$ as above. Without loss of generality, assume that \bar{x} is singleton, λ_m is in Negation Normal Form, and x only appears in the literals of the form $\ell < x$, $x < u$, and $x = e$, where ℓ , u , and e are x -free. Let $lits(\lambda)$ denote the literals of λ . The LW-method states that

$$\exists x \cdot \lambda_m(x) \equiv \left(\bigvee_{(x=e) \in lits(\lambda)} \lambda_m[e] \vee \bigvee_{(\ell < x) \in lits(\lambda)} \lambda_m[\ell + \epsilon] \vee \lambda_m[-\infty] \right) \quad (2)$$

where $\lambda_m[\cdot]$ denotes a *virtual substitution* for the literals containing x . Intuitively, $\lambda_m[e]$ covers the case when a literal ($x = e$) is true. Otherwise, the set of ℓ 's in the literals ($\ell < x$) identify intervals in which x can lie which are covered by the remaining substitutions. We omit the details of the substitution and instead illustrate it on an example. Let λ_m be $(x = e \wedge \phi_1) \vee (\ell < x \wedge x < u) \vee (x < u \wedge \phi_2)$, where $\ell, e, u, \phi_1, \phi_2$ are x -free. Then,

$$\begin{aligned} \exists x \cdot \lambda_m(x) &\equiv \lambda_m[e] \vee \lambda_m[\ell + \epsilon] \vee \lambda_m[-\infty] \\ &\equiv (\phi_1 \vee (\ell < e \wedge e < u) \vee (e < u \wedge \phi_2)) \vee (\ell < u \vee (\ell < u \wedge \phi_2)) \vee \phi_2 \\ &\equiv \phi_1 \vee (\ell < u) \vee \phi_2 \end{aligned}$$

We now define an MBP $LRAProj_\lambda$ for LRA as a map from models of λ_m to disjuncts in (2). Given $M \models \lambda_m$, $LRAProj_\lambda$ picks a disjunct that covers M based on values of the literals of the form $x = e$ and $\ell < x$ in M . Ties are broken by a syntactic ordering on terms (e.g., when $M \models \ell' = \ell$ for two literals $\ell < x$ and $\ell' < x$).

$$LRAProj_\lambda(M) = \begin{cases} \lambda_m[e], & \text{if } (x = e) \in lits(\lambda) \wedge M \models x = e \\ \lambda_m[\ell + \epsilon], & \text{else if } (\ell < x) \in lits(\lambda) \wedge M \models \ell < x \wedge \\ & \forall (\ell' < x) \in lits(\lambda) \cdot M \models ((\ell' < x) \Rightarrow (\ell' \leq \ell)) \\ \lambda_m[-\infty], & \text{otherwise} \end{cases}$$

Theorem 4. *LRAProj $_\lambda$ is a Model Based Projection.*

Note that $LRAProj_\lambda$ is linear in the size of λ . An MBP for LIA can be defined similarly (see Appendix) based on Cooper's method [15].

Bounded Safety with MBP. Intuitively, each quantifier-free formula in the image of $Proj_\lambda$ under-approximates λ . As above, we use λ_m for the quantifier-free matrix of λ . We modify the side-condition $\psi = \lambda$ of REACH and QUERY to use quantifier-free under-approximations as follows: (i) for REACH, the new side-condition is $\psi = Proj_\lambda(M)$ where $M \models \lambda_m \wedge \varphi$, (ii) for QUERY, the new side-condition is $\psi = Proj_\lambda(M)$ where $M \models \lambda_m \wedge \llbracket \Sigma_R(a) \rrbracket_\sigma^{b-1}$. Note that to avoid redundant applications of the rules, we require M to satisfy a formula stronger than λ_m . Intuitively, (i) ensures that the newly inferred reachability

	SLAM		SVCOMP-1		SVCOMP-2		SVCOMP-3	
	SAFE	UNSAFE	SAFE	UNSAFE	SAFE	UNSAFE	SAFE	UNSAFE
SPACER	1,721	985	249	509	213	497	234	482
Z3	1,722	997	245	509	208	493	234	477
VBS	1,727	998	252	509	225	500	240	482

Fig. 8: Number of programs verified by SPACER, Z3 and the Virtual Best Solver.

fact answers the current query and (ii) ensures that the new query cannot be immediately answered by known facts. In both cases, the required model M can be obtained as a side-effect of discharging the premises of the rules. Soundness of `BNDSAFETY` is unaffected and termination of `BNDSAFETY` follows from the image-finiteness of $Proj_\lambda$.

Theorem 5. *Assuming an oracle and an MBP for Th , `BNDSAFETY` is sound and terminating with the modified rules.*

Thus, `BNDSAFETY` with a linear-time MBP (such as $LRAProj_\lambda$) keeps the size of the formulas small by efficiently inferring only the necessary under-approximations of the quantified formulas.

6 Implementation and Experiments

We have implemented RECMC for analyzing C programs as part of the tool SPACER. The back-end is based on Z3 [18] which is used for SMT-solving and interpolation. It supports propositional logic, linear arithmetic, and bit-vectors (via bit-blasting). The front-end is based on UFO [3]. It converts C programs to the Horn-SMT format of Z3, which corresponds to the logical program representation of Section 3. The implementation and benchmarks are available online⁴.

We evaluated SPACER on two sets of benchmarks. The first set contains 2,908 Boolean Programs obtained from the SLAM toolkit⁵. The second contains 799 C programs from the Software Verification Competition (SVCOMP) 2014 [37]. We call this set SVCOMP-1. We also evaluated on two variants of SVCOMP-1, which we call SVCOMP-2 and SVCOMP-3, obtained by factoring out parts of the program into procedures and introducing more modularity. We compared SPACER against the implementation of GPDR in Z3. We used a time limit of 30 minutes and a memory limit of 16GB, on an Ubuntu machine with a 2.2 GHz AMD Opteron(TM) Processor 6174 and 516GB RAM. The results are summarized in Fig. 8. Since there are programs verified by only one of the tools, Fig. 8 also reports the number of programs verified by at least one, *i.e.*, the Virtual Best Solver (VBS).

Boolean Program Benchmarks. On most of the SLAM benchmarks, the runtimes of SPACER and Z3 are similar (within 2 minutes). We then evaluated on a Boolean Program from [6] in which the size of the call-tree grows exponentially in the number of procedures. As Fig. 9(a) shows, SPACER handles the increasing complexity in the example significantly better than Z3.

⁴ <http://www.cs.cmu.edu/~akomurav/projects/spacer/home.html>.

⁵ <https://svn.sosy-lab.org/software/sv-benchmarks/trunk/clauses/BOOL/slam.zip>

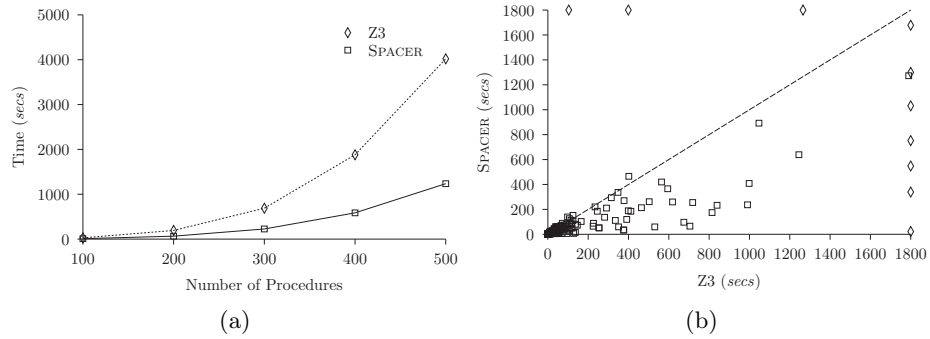


Fig. 9: SPACER vs. Z3 for (a) BEBOP example and (b) SVCOMP-1 benchmarks.




Fig. 10: SPACER vs. Z3 for the benchmarks (a) SVCOMP-2 and (b) SVCOMP-3.

SVCOMP 2014 Benchmarks. Fig. 9(b), 10(a) and 10(b) show the scatter plots for SVCOMP-1, SVCOMP-2 and SVCOMP-3 benchmarks. A diamond indicates a time-out and a star indicates a mem-out. The plots show that SPACER is significantly better on most of the programs. This shows the practical advantage of the approximations and MBP of RECMC.

7 Related Work

There is a large body of work on interprocedural program analysis. It was pointed out early on that verification of recursive programs is reducible to the computation of a fixed-point over relations (called *summaries*) representing the input-output behavior of each procedure [11,36]. Such procedure summaries are called *partial correctness relations* in [11], and are part of the *functional approach* of [36]. Reps, Horwitz, and Sagiv [35] showed that for a large class of finite interprocedural dataflow problems the summaries can be computed in time polynomial in the number of *facts* and procedures. Ball and Rajamani [6] adapted the RHS algorithm to the verification of Boolean Programs. Following the SLAM project, other software model checkers – such as BLAST [26] and MAGIC [10]

– also implemented the CEGAR loop with predicate abstraction. None used under-approximations of procedure semantics as we do.

Recently, several SMT-based algorithms have been proposed for safety verification of recursive programs, including WHALE [2], HSF [21], Duality [32], Ultimate Automizer [24], and Corral [29]. While these algorithms have been developed independently, they share a similar structure. They use SMT-solvers to look for counterexamples and interpolation to over-approximate summaries. Corral is an exception, which relies on user input and heuristics to supply the summaries. The algorithms differ in the SMT encoding and the heuristics used. However, in the worst-case, they completely unroll the call graph into a tree.

The work closest to ours is Generalized Property Driven Reachability (GPDR) of Hoder and Bjørner [27]. GPDR extends the hardware model checking algorithm IC3 of Bradley [9] to SMT-supported theories and recursive programs. Unlike RECMC, GPDR does not maintain reachability facts. In the context of Fig. 6, this means that ρ is always empty and there is no REACH rule. Instead, the QUERY rule is modified to use a model M that satisfies the premises (instead of our use of the path π when creating a query). Furthermore, the answers to the queries are cached. In the context of Boolean Programs, this ensures that every query is asked at most once (and either cached or blocked by a summary fact). Since there are only finitely many models, the algorithm always terminates. However, in the case of Linear Arithmetic, a formula can have infinitely many models and GPDR might end up applying the QUERY rule indefinitely. In contrast, RECMC creates only finitely many queries for a given bound on the call-stack and is guaranteed to find a counterexample if one exists.

Combination of over- and under-approximations for analysis of procedural programs has also been explored in [23,20]. However, our notion of an under-approximation is very different. Both [23,20] under-approximate summaries by *must transitions*. A must transition is a pair of formulas $\langle \varphi, \psi \rangle$ that under-approximates the summary of a procedure P iff for every state that satisfies φ , P has an execution that ends in a state satisfying ψ . In contrast, our reachability facts are similar to *summary edges* of RHS [35]. A reachability fact is a single formula φ such that every satisfying assignment to φ captures a terminating execution of P .

8 Conclusion

We presented RECMC, a new SMT-based algorithm for model checking safety properties of recursive programs. For programs and properties over decidable theories, RECMC is guaranteed to find a counterexample if one exists. To our knowledge, this is the first SMT-based algorithm with such a guarantee while being polynomial for Boolean Programs. The key idea is to use a combination of under- and over-approximations of the semantics of procedures, avoiding re-exploration of parts of the state-space. We described an efficient instantiation of RECMC for Linear Arithmetic (over rationals and integers) by introducing *Model-Based Projection* to under-approximate the expensive quantifier elimina-

tion. We have implemented it in our tool SPACER and shown empirical evidence that it significantly improves on the state-of-the-art.

In the future, we would like to explore extensions to other theories. Of particular interest are the theory EUF of uninterpreted functions with equality and the theory of arrays. The challenge is to deal with the lack of quantifier elimination. Another direction of interest is to combine RECMC with *Proof-based Abstraction* [31,22,28] to explore a combination of the approximations of procedure semantics with transition-relation abstraction.

Acknowledgment

We thank Edmund M. Clarke and Nikolaj Bjørner for many helpful discussions. Our definition of MBP is based on the idea of projected implicants co-developed with Nikolaj. We thank Cesare Tinelli and the anonymous reviewers for insightful comments. This research was sponsored by the National Science Foundation grants no. DMS1068829, CNS0926181 and CNS0931985, the GSRC under contract no. 1041377, the Semiconductor Research Corporation under contract no. 2005TJ1366, the Office of Naval Research under award no. N000141010188 and the CMU-Portugal Program. This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense. This material has been approved for public release and unlimited distribution. DM-0000973.

References

1. A. Albarghouthi, A. Gurfinkel, and M. Chechik. From Under-Approximations to Over-Approximations and Back. In *TACAS*, 2012.
2. A. Albarghouthi, A. Gurfinkel, and M. Chechik. Whale: An Interpolation-Based Algorithm for Inter-procedural Verification. In *VMCAI*, pages 39–55, 2012.
3. A. Albarghouthi, A. Gurfinkel, Y. Li, S. Chaki, and M. Chechik. UFO: Verification with Interpolants and Abstract Interpretation - (Competition Contribution). In *TACAS*, 2013.
4. R. Alur, M. Benedikt, K. Etessami, P. Godefroid, T. Reps, and M. Yannakakis. Analysis of Recursive State Machines. *TOPLAS*, 27(4):786–818, 2005.
5. T. Ball, R. Majumdar, T. Millstein, and S. K. Rajamani. Automatic Predicate Abstraction of C Programs. *SIGPLAN Not.*, 36(5):203–213, 2001.
6. T. Ball and S. K. Rajamani. Bebop: A Symbolic Model Checker for Boolean Programs. In *SPIN*, pages 113–130, 2000.
7. M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In *FMCO*, pages 364–387, 2005.
8. A. Biere, A. Cimatti, E. M. Clarke, O. Strichman, and Y. Zhu. Bounded Model Checking. *Advances in Computers*, 58:117–148, 2003.
9. A. R. Bradley. SAT-Based Model Checking without Unrolling. In *VMCAI*, 2011.
10. S. Chaki, E. M. Clarke, A. Groce, S. Jha, and H. Veith. Modular Verification of Software Components in C. *IEEE Trans. Software Eng.*, 30(6):388–402, 2004.
11. E. M. Clarke. Program Invariants as Fixed Points. *Computing*, 21(4):273–294, 1979.
12. E. M. Clarke. Programming Language Constructs for Which It Is Impossible To Obtain Good Hoare Axiom Systems. *JACM*, 26(1):129–147, 1979.
13. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-Guided Abstraction Refinement. In *CAV*, 2000.

14. E. M. Clarke, D. Kroening, and F. Lerda. A Tool for Checking ANSI-C Programs. In *TACAS*, 2004.
15. D. C. Cooper. Theorem Proving in Arithmetic without Multiplication. *Machine Intelligence*, pages 91–100, 1972.
16. W. Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *Symbolic Logic*, 22(3):269–285, 1957.
17. L. De Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *TACAS*, 2008.
18. L. M. de Moura and N. Bjørner. Z3: An Efficient SMT Solver. In C. R. Ramakrishnan and J. Rehof, editors, *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 337–340. Springer, 2008.
19. J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. Efficient Algorithms for Model Checking Pushdown Systems. In *CAV*, CAV '00, pages 232–247, 2000.
20. P. Godefroid, A. V. Nori, S. K. Rajamani, and S. Tetali. Compositional May-Must Program Analysis: Unleashing the Power of Alternation. In *POPL*, pages 43–56, 2010.
21. S. Grebenschikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing Software Verifiers from Proof Rules. In *PLDI*, pages 405–416, 2012.
22. A. Gupta, M. K. Ganai, Z. Yang, and P. Ashar. Iterative Abstraction using SAT-based BMC with Proof Analysis. In *ICCAD*, pages 416–423, 2003.
23. A. Gurfinkel, O. Wei, and M. Chechik. Model checking recursive programs with exact predicate abstraction. In S. D. Cha, J.-Y. Choi, M. Kim, I. Lee, and M. Viswanathan, editors, *ATVA*, volume 5311 of *Lecture Notes in Computer Science*, pages 95–110. Springer, 2008.
24. M. Heizmann, J. Christ, D. Dietsch, E. Ermis, J. Hoenicke, M. Lindenmann, A. Nutz, C. Schilling, and A. Podelski. Ultimate Automizer with SMTInterpol - (Competition Contribution). In N. Piterman and S. A. Smolka, editors, *TACAS*, volume 7795 of *Lecture Notes in Computer Science*, pages 641–643. Springer, 2013.
25. M. Heizmann, J. Hoenicke, and A. Podelski. Nested Interpolants. *SIGPLAN Not.*, 45(1)(1):471–482, 2010.
26. T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy abstraction. In *Proc. of POPL*, pages 58–70, 2002.
27. K. Hoder and N. Bjørner. Generalized Property Directed Reachability. In *SAT*, 2012.
28. A. Komuravelli, A. Gurfinkel, S. Chaki, and E. M. Clarke. Automated Abstraction in SMT-based Unbounded Software Model Checking. In *CAV*, pages 846–862, 2013.
29. A. Lal, S. Qadeer, and S. K. Lahiri. A solver for reachability modulo theories. In P. Madhusudan and S. A. Seshia, editors, *CAV*, volume 7358 of *Lecture Notes in Computer Science*, pages 427–443. Springer, 2012.
30. R. Loos and V. Weispfenning. Applying Linear Quantifier Elimination. *Computing*, 36(5):450–462, 1993.
31. K. L. McMillan and N. Amla. Automatic Abstraction without Counterexamples. In *TACAS*, 2003.
32. K. L. McMillan and A. Rybalchenko. Solving Constrained Horn Clauses using Interpolation. Technical Report MSR-TR-2013-6, Microsoft Research, 2013.
33. R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(*T*). *J. ACM*, 53(6):937–977, 2006.
34. T. Nipkow. Linear Quantifier Elimination. *J. Autom. Reason.*, 45(2):189–212, 2010.

35. T. W. Reps, S. Horwitz, and S. Sagiv. Precise Interprocedural Dataflow Analysis via Graph Reachability. In *POPL*, pages 49–61, 1995.
36. M. Sharir and A. Pnueli. *Program Flow Analysis: Theory and Applications*, chapter “Two Approaches to Interprocedural Data Flow Analysis”, pages 189–233. Prentice-Hall, 1981.
37. Software Verification Competition. TACAS, 2014.
<http://sv-comp.sosy-lab.org>.

A Divergence of GPDR for Bounded Call-Stack

Consider the program $\langle\langle M, L, G \rangle, M\rangle$ with $M = \langle y_0, y, \Sigma_M, \langle x, n \rangle, \beta_M \rangle$, $L = \langle n, \langle x, y, i \rangle, \Sigma_L, \langle x_0, y_0, i_0 \rangle, \beta_L \rangle$, $G = \langle x_0, x_1, \Sigma_G, \emptyset, \beta_G \rangle$, where

$$\begin{aligned}\beta_M &= \Sigma_L(x, y_0, n, n) \wedge \Sigma_G(x, y) \wedge n > 0 \\ \beta_L &= (i = 0 \wedge x = 0 \wedge y = 0) \vee \\ &\quad (\Sigma_L(x_0, y_0, i_0, n) \wedge x = x_0 + 1 \wedge y = y_0 + 1 \wedge i = i_0 + 1 \wedge i > 0) \\ \beta_G &= (x = x_0 + 1)\end{aligned}$$

The GPDR [27] algorithm can be shown to diverge when checking $M \models_2 y_0 \leq y$, for e.g., by inferring the diverging sequence of over-approximations of $\llbracket L \rrbracket^1$:

$$(x < 2 \Rightarrow y \leq 1), (x < 3 \Rightarrow y \leq 2), \dots$$

We also observed this behavior experimentally (Z3 revision d548c51 at <http://z3.codeplex.com>). The Horn-SMT file for the example is available at

http://www.cs.cmu.edu/~akomurav/projects/spacer/gpdr_diverging.smt2.

B Soundness of RECMC and BNDSAFETY (Proof of Theorem 1)

We first restate the theorem.

Theorem 1. *RECMC and BNDSAFETY are sound.*

Proof. We only show the soundness of BNDSAFETY; the soundness of RECMC easily follows. In particular, for $\text{BNDSAFETY}(M, \varphi_{\text{safe}}, n, \emptyset, \emptyset)$ we show the following:

1. if the premises of UNSAFE hold, $M \not\models_n \varphi_{\text{safe}}$, and
2. if the premises of SAFE hold, $M \models_n \varphi_{\text{safe}}$.

It suffices to show that the environments U_ρ^b and O_σ^b , respectively, under- and over-approximate the bounded semantics of the procedures, for every $0 \leq b \leq n$. In particular, we show that the following is an invariant of BNDSAFETY: for every model I of the background theory Th , for every $Q \in \mathcal{A}$ and $b \in [0, n]$,

$$I(U_\rho^b(\Sigma_Q)) \subseteq \llbracket Q \rrbracket_I^b \subseteq I(O_\sigma^b(\Sigma_Q)). \quad (3)$$

Initially, ρ and σ are empty and the invariant holds trivially. BNDSAFETY updates σ and ρ in the rules SUM and REACH, respectively. We show that these rules preserve (3). We only show the case of SUM. The case of REACH is similar.

Let $\langle P, \varphi, b \rangle \in \mathcal{Q}$ be such that SUM is applicable, *i.e.*, $\models \llbracket \beta_P \rrbracket_{\sigma}^{b-1} \Rightarrow \neg \varphi$. Let $\psi = \text{ITP}(\llbracket \beta_P \rrbracket_{\sigma}^{b-1}, \neg \varphi)$. Note that φ , and hence ψ , does not depend on the local variables ℓ_P . Hence, we know that

$$\models (\exists \bar{\ell}_P \cdot \llbracket \beta_P \rrbracket_{\sigma}^{b-1}) \Rightarrow \psi. \quad (4)$$

The case of $b = 0$ is easy and we will skip it. Let I be an arbitrary model of Th . Assume that (3) holds at $b - 1$ before applying the rule. In particular, assume that for all $Q \in \mathcal{A}$, $\llbracket Q \rrbracket_I^{b-1} \subseteq I(O_{\sigma}^{b-1}(\Sigma_Q))$.

We will first show that the new summary fact ψ over-approximates $\llbracket P \rrbracket_I^b$. Let $J(I, \bar{X})$ be an expansion of I as defined in Section 3.

$$\begin{aligned} \llbracket P \rrbracket_I^b &= J(I, \langle \llbracket P_0 \rrbracket_I^{b-1}, \dots, \llbracket P_n \rrbracket_I^{b-1} \rangle \rangle (\exists \bar{\ell}_{P_i} \cdot \beta_{P_i})) \\ &= J(I, \langle I(O_{\sigma}^{b-1}(\Sigma_{P_0})), \dots, I(O_{\sigma}^{b-1}(\Sigma_{P_n})) \rangle \rangle (\exists \bar{\ell}_{P_i} \cdot \beta_{P_i})) && \text{(hypothesis)} \\ &= I(\llbracket \exists \bar{\ell}_P \cdot \beta_P \rrbracket_{O_{\sigma}^{b-1}}) && (O_{\sigma}^{b-1} \text{ is FO-definable}) \\ &= I(\exists \bar{\ell}_P \cdot \llbracket \beta_P \rrbracket_{O_{\sigma}^{b-1}}) && \text{(logic)} \\ &= I(\exists \bar{\ell}_P \cdot \llbracket \beta_P \rrbracket_{\sigma}^{b-1}) && \text{(notation)} \\ &\subseteq I(\psi) && \text{(from (4))} \end{aligned}$$

Next, we show that the invariant continues to hold. The map of summary facts is updated to $\sigma' = \sigma \cup \{\langle P, b \rangle \mapsto \psi\}$. Now, σ' differs from σ only for the procedure P and every bound in $[0, b]$. Let $b' \in [0, b]$ be arbitrary. Since (3) was true before applying SUM, we know that $\llbracket P \rrbracket_I^{b'} \subseteq I(O_{\sigma}^{b'}(\Sigma_P))$. As $\llbracket P \rrbracket_I^{b'} \subseteq \llbracket P \rrbracket_I^b \subseteq I(\psi)$, it follows that $\llbracket P \rrbracket_I^{b'} \subseteq I(O_{\sigma}^{b'}(\Sigma_P)) \cap I(\psi) \subseteq I(O_{\sigma'}^{b'}(\Sigma_P) \wedge \psi) = I(O_{\sigma'}^{b'}(\Sigma_P))$. \square

C Termination of BNDSAFETY (Proof of Theorem 2)

We first restate the theorem:

Theorem 2. *Given an oracle for Th , $\text{BNDSAFETY}(\mathcal{A}, \varphi, n, \emptyset, \emptyset)$ terminates.*

We begin with showing some useful lemmas. Let p be the maximum number of paths in $P \in \mathcal{A}$, c be the maximum number of procedure calls along any path in \mathcal{A} , N be the number of procedures in \mathcal{A} and assume an oracle for SAT modulo Th .

The following lemma shows that when a query is removed from \mathcal{Q} , it is actually answered. The proof is immediate from the definitions of O_{σ}^b and U_{ρ}^b given in Section 3.

Lemma 1 (Answered Queries). *Whenever BNDSAFETY removes a query from \mathcal{Q} , it is answered using the known summary and reachability facts. In particular, let SUM or REACH be applied to $\langle P, \varphi, b \rangle$. Then, for every $\langle P, \eta, b \rangle \in \mathcal{Q}$ removed from \mathcal{Q} by the rule,*

1. *after SUM is applied, $\models \llbracket \Sigma_P \rrbracket_{\sigma}^b \Rightarrow \neg \eta$, and*

2. after REACH is applied, $\not\models \llbracket \Sigma_P \rrbracket_\rho^b \Rightarrow \neg\eta$.

Next, we show that new facts need to be inferred to answer queries remaining in \mathcal{Q} .

Lemma 2 (Pending Queries). *\mathcal{Q} only has the queries which cannot be immediately answered by ρ or σ , i.e., as long as $\langle P, \eta, \ell \rangle$ is in \mathcal{Q} , the following are invariant across iterations of BND SAFETY.*

1. $\not\models \llbracket \Sigma_P \rrbracket_\sigma^\ell \Rightarrow \neg\eta$, and
2. $\models \llbracket \Sigma_P \rrbracket_\rho^\ell \Rightarrow \neg\eta$.

Proof. We first show that the invariants hold when a query is newly created by QUERY. Let P , η and ℓ be, respectively, R , $\psi[\bar{a} \leftarrow \bar{v}_R]$ and $b - 1$, as in the conclusion of the rule. The last-but-one premise of QUERY is

$$\models \llbracket \pi_u \rrbracket_\sigma^{b-1} \wedge \llbracket \Sigma_R(\bar{a}) \rrbracket_\rho^{b-1} \wedge \llbracket \pi_v \rrbracket_\rho^{b-1} \Rightarrow \neg\varphi$$

which implies that

$$\models \llbracket \Sigma_R(\bar{a}) \rrbracket_\rho^{b-1} \Rightarrow \neg (\llbracket \pi_u \rrbracket_\sigma^{b-1} \wedge \llbracket \pi_v \rrbracket_\rho^{b-1} \wedge \varphi).$$

The variables not in common, *viz.*, $(\bar{v}_P \cup \bar{\ell}_P) \setminus \bar{a}$, can be universally quantified from the right hand side resulting in $\models \llbracket \Sigma_R \rrbracket_\rho^{b-1} \Rightarrow \neg\eta$. Similarly, $\not\models \llbracket \Sigma_R \rrbracket_\sigma^{b-1} \Rightarrow \neg\eta$ follows from the last premise of the rule. Next, we show that SUM and REACH preserve the invariants.

Let SUM answer a query $\langle P, \varphi, \ell \rangle$ with a new summary fact ψ and let the updated map of summary facts be $\sigma' = \sigma \cup \{\langle P, \ell \rangle \mapsto \psi\}$. Now, consider $\langle P, \eta, \ell' \rangle \in \mathcal{Q}$ after the application of the rule. If $\ell' > \ell$, $O_{\sigma'}^{\ell'} = O_\sigma^{\ell'}$ and the invariant continues to hold. So, assume $\ell' \leq \ell$. From the conclusion of SUM, we have $\not\models \llbracket \Sigma_P \rrbracket_\sigma^{\ell'} \wedge \psi \Rightarrow \neg\eta$. Now, $\llbracket \Sigma_P \rrbracket_{\sigma'}^{\ell'} = \llbracket \Sigma_P \rrbracket_\sigma^{\ell'} \wedge \psi$. So, the invariant continues to hold.

Similarly, let REACH answer a query $\langle P, \varphi, \ell \rangle$ with a new reachability fact ψ and let the updated map of reachability facts be $\rho' = \rho \cup \{\psi \mapsto \langle P, \ell \rangle\}$. Now, consider $\langle P, \eta, \ell' \rangle \in \mathcal{Q}$ after the application of the rule. If $\ell' < \ell$, $U_{\rho'}^{\ell'} = U_\rho^{\ell'}$ and the invariant continues to hold. So, assume $\ell' \geq \ell$. From the conclusion of REACH, we have $\models \psi \Rightarrow \neg\eta$. Assuming the invariant holds before the rule application, we also have $\models \llbracket \Sigma_P \rrbracket_\rho^{\ell'} \Rightarrow \neg\eta$. Therefore, we have $\models \llbracket \Sigma_P \rrbracket_\rho^{\ell'} \vee \psi \Rightarrow \neg\eta$. Now, $\llbracket \Sigma_P \rrbracket_{\rho'}^{\ell'} = \llbracket \Sigma_P \rrbracket_\rho^{\ell'} \vee \psi$. So, the invariant continues to hold. \square

The next few lemmas show that the rules of the algorithm cannot be applied indefinitely, leading to a termination argument. Let N be the number of procedures in the program, p be the maximum number of paths in a procedure, c be the maximum number of procedure calls along any path in \mathcal{A} and n be the current bound on the call-stack.

Lemma 3 (Finite Reach Facts). *The environment U_ρ^b is updated for a given predicate symbol Σ_P and a bound b on the call-stack only $O(N^b \cdot p^{b+1})$ -many times.*

Proof. The environment U_ρ^b can be updated for Σ_P and b whenever a reachability fact is inferred for P at a bound $b' \leq b$. Now, a reachability fact is obtained per path (after eliminating the local variables) of a procedure, using the currently known reachability facts about the callees. Moreover, Lemmas 1 and 2 imply that no reachability fact is inferred twice. This is because whenever a query is answered using REACH, the query could not have been answered using already existing reachability facts and a new reachability fact is inferred.

This gives the following recurrence $Reach(b)$ for the number of updates to U_ρ^b for a given Σ_P :

$$Reach(b) = \begin{cases} p, & b = 0 \\ (p \cdot N + 1) \cdot Reach(b-1), & b > 0. \end{cases}$$

In words, for $b = 0$, the number of updates is given by the number of reachability facts that can be inferred, which is bounded by the number of paths p in the procedure P . For $b > 0$, the environment U_ρ^{b-1} is updated when a reachability fact is learnt at b or at a bound smaller than b . The latter is simply $Reach(b-1)$. For the former, a new reachability fact is inferred at b along a path whenever U_ρ^{b-1} changes for a callee. For N procedures and p paths, this is given by $(p \cdot N \cdot Reach(b-1))$.

This gives us $Reach(b) = O(N^b \cdot p^{b+1})$. □

Lemma 4 (Finite Queries). *For $\langle P, \varphi, b \rangle \in \mathcal{Q}$, QUERY is applicable only $O(c \cdot N^b \cdot p^{b+1})$ -many times.*

Proof. First, assume that the environments U_ρ^{b-1} and O_σ^{b-1} are fixed. The number of possible queries that can be created for a given path of P is bounded by the number of ways the path can be divided into a prefix, a procedure call, and a suffix. This is bounded by c , the maximum number of calls along the path. For p paths, this is bounded by $c \cdot p$.

Consider a path π and its division, and let a query be created for a callee R along π . Now, while the query is still in \mathcal{Q} , updates to the environments O_σ^{b-1} and U_ρ^{b-1} do not result in a new query for R for the same division along π . This is because, the new query would overlap with the existing one and this is disallowed by the second side-condition of QUERY.

Suppose that the new query is answered by SUM. With the updated map of summary facts, the last premise of QUERY can be shown to fail for the current division of π . If O_σ^{b-1} is updated, the last premise continues to fail. So, a new query can be created for the same prefix and suffix along π only if U_ρ^{b-1} is updated for some callee along π . The other possibility is that the query is answered by REACH which updates U_ρ^{b-1} as well.

Thus, for a given path, and a given division of it into prefix and suffix, the number of queries that can be created is bounded by the number of updates to U_ρ^{b-1} which is $(N \cdot Reach(b-1))$. Here, $Reach$ is as in Lemma 3. So, the number of times QUERY is applicable for a given query $\langle P, \varphi, b \rangle$ is $O(p \cdot c \cdot N \cdot Reach(b-1))$. As $Reach(b) = N^b \cdot p^{b+1}$, we obtain the bound $O(c \cdot N^b \cdot p^{b+1})$. □

Lemma 5 (Progress). *As long as \mathcal{Q} is non-empty in BND SAFETY, either SUM, REACH or QUERY is always applicable.*

Proof. First, we show that for every query in \mathcal{Q} , either of the three rules is applicable, without the second side-condition in QUERY. Let $\langle P, \varphi, b \rangle \in \mathcal{Q}$. If $\models \llbracket \beta_P \rrbracket_\sigma^{b-1} \Rightarrow \neg\varphi$, then SUM is applicable. Otherwise, there exists a path $\pi \in \text{Paths}(P)$ such that $\llbracket \pi \rrbracket_\sigma^{b-1}$ is satisfiable with φ , i.e., $\not\models \llbracket \pi \rrbracket_\sigma^{b-1} \Rightarrow \neg\varphi$. Now, if $\llbracket \pi \rrbracket_\rho^{b-1}$ is also satisfiable with φ , i.e., $\not\models \llbracket \pi \rrbracket_\rho^{b-1} \Rightarrow \neg\varphi$, REACH is applicable. Otherwise, $\models \llbracket \pi \rrbracket_\rho^{b-1} \Rightarrow \neg\varphi$. Note that this can only happen if $b > 0$, as otherwise, there will not be any procedure calls along π and $\llbracket \pi \rrbracket_\sigma^{b-1}$ and $\llbracket \pi \rrbracket_\rho^{b-1}$ would be equivalent.

Let $\pi = \pi_0 \wedge \pi_1 \wedge \dots \wedge \pi_l$ for some finite l . Then, $\llbracket \pi \rrbracket_\sigma^{b-1}$ is obtained by taking the conjunction of the formulas

$$\langle \llbracket \pi_0 \rrbracket_\sigma^{b-1}, \llbracket \pi_1 \rrbracket_\sigma^{b-1}, \dots \rangle.$$

Similarly, $\llbracket \pi \rrbracket_\rho^{b-1}$ is obtained by taking the conjunction of the formulas

$$\langle \llbracket \pi_0 \rrbracket_\rho^{b-1}, \llbracket \pi_1 \rrbracket_\rho^{b-1}, \dots \rangle.$$

From Theorem 1, we can think of obtaining the latter sequence of formulas by conjoining $\llbracket \pi_i \rrbracket_\rho^{b-1}$ to $\llbracket \pi_i \rrbracket_\sigma^{b-1}$ for every i . When this is done backwards for decreasing values of i , an intermediate sequence looks like

$$\langle \llbracket \pi_0 \rrbracket_\sigma^{b-1}, \dots, \llbracket \pi_{j-1} \rrbracket_\sigma^{b-1}, \llbracket \pi_j \rrbracket_\rho^{b-1} \dots \rangle.$$

As $\llbracket \pi \rrbracket_\rho^{b-1}$ is unsatisfiable with φ , there exists a maximal j such that the conjunction of constraints in such an intermediate sequence are unsatisfiable with φ . Moreover, π_j must be a literal of the form $\Sigma_R(\bar{a})$ as otherwise, $\llbracket \pi_j \rrbracket_\sigma^{b-1} = \llbracket \pi_j \rrbracket_\rho^{b-1}$ violating the maximality condition on j . Thus, all premises of QUERY hold and the rule is applicable.

Now, the second side-condition in QUERY can be trivially satisfied by always choosing a query in \mathcal{Q} with the smallest bound for the next rule to apply. This is because, if $\langle R, \eta, b-1 \rangle$ is the newly created query, there is no other query in \mathcal{Q} for R and $b-1$. \square

Lemmas 4 and 5 imply that every query in \mathcal{Q} is eventually answered by SUM or REACH, as shown below.

Lemma 6 (Eventual Answer). *Every $\langle P, \varphi, b \rangle \in \mathcal{Q}$ is eventually answered by SUM or REACH, in $O(b \cdot c^b \cdot (Np)^{O(b^2)})$ applications of the rules.*

Proof. Firstly, to answer any given query in \mathcal{Q} , Lemma 4 guarantees that the algorithm can only create finitely many queries. Lemma 5 guarantees that some rule is always applicable, as long as \mathcal{Q} is non-empty. Thus, when QUERY cannot be applied for any query in \mathcal{Q} , either SUM or REACH must be applicable for some query. Thus, eventually, all queries are answered.

The total number of rule applications to answer $\langle P, \varphi, b \rangle$ is then linear in the cumulative number of applications of QUERY, which has the following recurrence:

$$T(b) = \begin{cases} Q(0), & b = 0 \\ Q(b)(1 + T(b-1)), & b > 0. \end{cases}$$

where $Q(b)$ denotes the number of applications of QUERY for a fixed query in \mathcal{Q} at bound b . From Lemma 4, $Q(b) = O(c \cdot N^b \cdot p^{b+1})$. This gives us $T(b) = O(b \cdot c^b \cdot (Np)^{O(b^2)})$. \square

Main Proof. Follows immediately from Lemma 6. \square

D Complexity of RECMC for Boolean Programs (Proof of Theorem 3)

We first restate the theorem:

Theorem 3. *Let \mathcal{A} be a Boolean Program. Then RECMC(\mathcal{A}, φ) terminates in $O(N^2 \cdot 2^{2k})$ -many applications of the rules in Fig. 6.*

Proof. First, assume a bound n on the call-stack. The number of queries that can be created for a procedure at any given bound is $O(2^k)$, the number of possible valuations of the parameters (note that QUERY disallows overlapping queries to be present simultaneously in \mathcal{Q}). For N procedures and n possible values of the bound, the complexity of BNDSAFETY($\mathcal{A}, \varphi, n, \emptyset, \emptyset$), for a Boolean Program, is $O(N \cdot 2^k \cdot n)$.

Now, the total number of summary facts that can be inferred for a procedure is also bounded by $O(2^k)$. As O_σ^b is monotonic in b , the number of iterations of RECMC is bounded by $O(N \cdot 2^k)$, the cumulative number of states of all procedures. Thus, we obtain the complexity of RECMC as $O(N^2 \cdot 2^{2k})$. \square

E BNDSAFETY with MBP (Proof of Theorem 5)

Here, we show that BNDSAFETY with MBP is sound and terminating.

First of all, in presence of MBP, SUM is unaffected and a reachability fact inferred by REACH is only strengthened. Thus, soundness of BNDSAFETY (Theorem 1) is preserved.

Then, it is easy to show that the modified side-conditions to REACH and QUERY preserve Lemmas 1 and 2 and we skip the proof.

Then, we will show that the *finite-image* property of an MBP preserves the finiteness of the number of reachability facts inferred and the number of queries generated by the algorithm. Let d be the size of the image of an MBP. In the proof of Lemma 3, the recurrence relation has an extra factor of d . The rest of the proof of finiteness of the number of reachability facts remains the same. Similarly, in the proof of Lemma 4, the number of times QUERY can be applied

along a path for a fixed division and fixed environments O_σ^{b-1} and U_ρ^{b-1} increases by a factor of d . Again, the rest of the proof of finiteness of the number of queries generated remains the same. That is, Lemmas 3 and 4, and hence, Lemma 6, are preserved with scaled up complexity bounds.

Note that Theorem 5 is unaffected by under-approximations.

Together, we have that Theorem 2 is preserved, with a scaled up complexity bound. \square

F $LRAProj_\lambda$ is an MBP (Proof of Lemma 4)

First, we restate the theorem:

Theorem 4. *$LRAProj_\lambda$ is a Model Based Projection.*

Proof. By definition, $LRAProj_\lambda$ has a finite image, as there are only finitely many disjuncts in (2). Thus, it suffices to show that for every $M \models \lambda_m$, $M \models LRAProj_\lambda(M)$.

Each disjunct in the LW decomposition (2) is obtained by a *virtual substitution* of the literals in λ_m containing x . As in Section 5, we assume that λ_m is in NNF with the only literals containing x of the form $(x = e)$, $(\ell < x)$ or $(x < u)$ for x -free terms e , ℓ and u . Let Sub_t denote the virtual substitution map of literals when t is either e , $\ell + \epsilon$ or $-\infty$. The LW method [30] defines:

$$Sub_e(x = e) = \top, Sub_e(\ell < x) = (\ell < e), Sub_e(x < u) = (e < u) \quad (5)$$

$$Sub_{\ell+\epsilon}(x = e) = \perp, Sub_{\ell+\epsilon}(\ell' < x) = (\ell' \leq \ell), Sub_{\ell+\epsilon}(x < u) = (\ell < u) \quad (6)$$

$$Sub_{-\infty}(x = e) = \perp, Sub_{-\infty}(\ell < x) = \perp, Sub_{-\infty}(x < u) = \top \quad (7)$$

Let $M \models \lambda_m$ and $LRAProj_\lambda(M) = \lambda_m[t]$ where t is either e or $\ell + \epsilon$ or $-\infty$. As λ_m is in NNF, it suffices to show that for every literal μ of λ_m containing x , the following holds:

$$M \models (\mu \Rightarrow Sub_t(\mu)) \quad (8)$$

We consider the different possibilities of t below. For a term η , let $M[\eta]$ denote the value of η in M .

Case $t = e$. In this case, we know that $M \models x = e$. Now, for a literal $\ell < x$,

$$\begin{aligned} M[\ell < x] &\Rightarrow M[\ell] < M[x] \\ &= M[\ell] < M[e] \\ &= M[\ell < e] \\ &= M[Sub_t(\ell < x)] \quad \{Sub_t(\ell < x) = (\ell < e)\}. \end{aligned}$$

Similarly, literals of the form $x < u$ and $x = e'$ can be considered.

Case $t = \ell + \epsilon$. In this case, we know that $M[\ell < x]$ is true, *i.e.*, $M[\ell] < M[x]$ and whenever $M[\ell' < x]$ is true, $M[\ell' \leq \ell]$ is also true. Now, for a literal $\ell' < x$,

$$\begin{aligned} M[\ell' < x] &\Rightarrow M[\ell' \leq \ell] \\ &= M[Sub_t(\ell' < x)] \quad \{Sub_t(\ell' < x) = (\ell' \leq \ell)\}. \end{aligned}$$

For a literal $x < u$,

$$\begin{aligned} M[x < u] &\Rightarrow M[x] < M[u] \\ &\Rightarrow M[\ell] < M[u] \quad \{M[\ell] < M[x]\} \\ &\Rightarrow M[\ell < u] \\ &= M[Sub_t(x < u)] \quad \{Sub_t(x < u) = (\ell < u)\} \end{aligned}$$

For a literal $x = e$, (8) vacuously holds as $M[x = e]$ is false.

Case $t = -\infty$. In this case, we know that $M[x = e]$ and $M[\ell < x]$ are false for every literal of the form $x = e$ and $\ell < x$. So, for such literals (8) vacuously holds. For a literal $x < u$, $Sub_t(x < u) = \top$ and hence, (8) holds again. \square

G Model Based Projection for Linear Integer Arithmetic

In this section, we present our MBP $LIAProj_\lambda$ for LIA. It is based on Cooper's method for Quantifier Elimination procedure [15]. Let $\lambda(\bar{y}) = \exists x \cdot \lambda_m(x, \bar{y})$, where λ_m is quantifier free and in negation normal form. Without loss of generality, let the only literals containing x be the form $\ell < x$, $x < u$, $x = e$ or $(d \mid \pm x + w)$, where $a \mid b$ denotes that a divides b , the terms ℓ , u , e and w are x -free, and $d \in \mathbb{Z} \setminus \{0\}$. Let $E = \{e \mid (x = e) \in lits(\lambda_m)\}$ be the set of equality terms of x and $L = \{\ell \mid (\ell < x) \in lits(\lambda_m)\}$ be the set of lower-bounds of x . Then, by Cooper's method,

$$\exists x \cdot \lambda_m(x, \bar{y}) \equiv \bigvee_{(x=e) \in lits(\lambda)} \lambda_m[e] \vee \bigvee_{(\ell < x) \in lits(\lambda)} \left(\bigvee_{i=0}^{D-1} \lambda_m[\ell + 1 + i] \right) \vee \bigvee_{i=0}^{D-1} \lambda_m^{-\infty}[i]. \quad (9)$$

where D is the least common multiple of all the divisors in the divisibility literals of λ_m , $[\cdot]$ denotes a substitution for x and $\lambda_m^{-\infty}$ is obtained from λ_m by substituting all non-divisibility literals as follows:

$$(\ell < x) \mapsto \perp \quad (x < u) \mapsto \top \quad (x = e) \mapsto \perp \quad (10)$$

Intuitively, the disjunction partitions the space of the possible values of x . A disjunct for $(x = e)$ covers the case when x is equal to an equality term. Otherwise, the lower-bounds identify various intervals in which x can be present. The disjuncts for $(\ell < x)$ cover the case when x satisfies a lower-bound, and the last disjunct is for the case when x is smaller than all lower-bounds. The

disjunction over the possible values of i covers the different ways in which the divisibility literals can be satisfied.

Model-based projection $LIAProj_\lambda$ is defined as follows, conflicts are resolved by some arbitrary, but fixed, syntactic ordering on terms:

$$LIAProj_\lambda(M) = \begin{cases} \lambda_m[e], & \text{if } x = e \in lits(\lambda) \wedge M \models (x = e) \\ \lambda_m[\ell + 1 + i_\ell], & \text{else if } (\ell < x) \in lits(\lambda) \wedge M \models (\ell < x) \wedge \\ & \forall (\ell' < x) \in lits(\lambda) \cdot (M \models ((\ell' < x) \Rightarrow (\ell' \leq \ell))) \\ \lambda_m^{-\infty}[i_{-\infty}], & \text{otherwise} \end{cases} \quad (11)$$

where $i_\ell = M[x - (\ell + 1)] \bmod D$, $i_{-\infty} = M[x] \bmod D$, and $M[x]$ is the value of x in M . The following lemma shows that $LIAProj_\lambda$ is indeed a model based projection. The proof is similar to that of $LRAProj_\lambda$.

Lemma 7. *LIAProj $_\lambda$ is a Model-Based Projection.*