

Maximal Specification Synthesis

Aws Albarghouthi

University of Wisconsin–Madison

Isil Dillig

University of Texas at Austin

Arie Gurfinkel *

Carnegie Mellon University

Abstract

Many problems in program analysis, verification, and synthesis require inferring specifications of unknown procedures. Motivated by a broad range of applications, we formulate the problem of *maximal specification inference*: Given a postcondition φ and a program P calling a set of unknown procedures $F_1, \dots, F_n$, what are the most permissive specifications of procedures F_i that ensure correctness of P ? In other words, we are looking for the smallest number of assumptions we need to make about the behaviours of F_i in order to prove that P satisfies its postcondition.

To solve this problem, we present a novel approach that utilizes a counterexample-guided inductive synthesis loop and reduces the maximal specification inference problem to *multi-abduction*. We formulate the novel notion of *multi-abduction* as a generalization of classical logical abduction and present an algorithm for solving multi-abduction problems. On the practical side, we evaluate our specification inference technique on a range of benchmarks and demonstrate its ability to synthesize specifications of kernel routines invoked by device drivers.

Categories and Subject Descriptors F.3.1 [Specifying and Verifying and Reasoning about Programs]: Specification techniques

Keywords verification; specification; synthesis

1. Introduction

In the world of program analysis, verification, and synthesis, many problems require inferring specifications of unknown procedures. Specifically, given a program P calling procedures $F_1, \dots, F_n$, we often want to answer the following question:

What is the weakest specification φ_i for each callee F_i that ensures correctness of program P ?

In other words, what are the minimal assumptions that we can make about the callees of P and still be able to prove P correct? Here, we view specifications as logical relations over a procedure’s parameters and returns; therefore, weakest (maximal) relations imply most permissive specifications. This question, which we refer to as

the *maximal specification inference* problem, appears in many different guises. We now survey a broad spectrum of problems that can be viewed through the lens of maximal specification inference:

- **Verification of open programs:** Many programs are written against libraries whose source code is unavailable (e.g., proprietary), too large and complex to be verified, or written in a different programming language. To verify such programs, we need to know specifications of these library methods.
In recent years, there has been a flurry of interest in *automatically* inferring weakest specifications of library code [7, 12, 48, 54]. For instance, in the context of taint-flow analysis for Android applications, Bastani et al. [12] study synthesizing taint-flow specifications of Android library methods from their usage contexts. They advocate presenting these specifications to a “human auditor for validation,” with the goal of avoiding the prohibitively expensive analysis of the whole Android ecosystem or the error-prone task of manually writing stubs.
- **Compositional interprocedural verification:** Maximal specification inference facilitates *top-down* compositional verification of programs. Starting from the main procedure, we can infer weakest necessary specifications of callees, which can then be verified recursively down the call graph in a modular fashion.
- **Modular program synthesis:** Maximal specification inference can also enable modular program synthesis. For instance, given an incomplete program *sketch* with holes (i.e., missing expressions), we can model each hole as a call to an unknown procedure. Once we infer a maximal specification for each hole, we can then use off-the-shelf code synthesis techniques [8, 34, 38, 51] to independently synthesize each unknown expression. Furthermore, if the goal is to fill holes with constant parameters [51], maximal specification inference can logically and succinctly characterize a large or even infinite solution space of possible parameter combinations.
- **Input-filter synthesis:** Input-filter generation—a key problem in security—involves synthesizing a procedure that rejects inputs that crash the program or violate certain correctness criteria (see, e.g., [21, 26, 42]). We can cast this problem as maximal specification inference by inserting an unknown procedure at the beginning of the program that reads the input and accepts it or rejects it. By synthesizing a maximal specification for the unknown procedure, we are effectively synthesizing the most permissive input filter—the one that accepts all good inputs and rejects all bad ones.
- **Infinite-state games and program repair:** Infinite-state games can be modeled as programs with unknown procedures representing possible player strategies (moves) [14, 15]. Maximal specification inference allows us to determine whether there exists a winning strategy for a player, and, if so, what that strategy is. In addition, following the observations of Jobstmann et al. [37], we can reduce program repair to a game-solving prob-

* Copyright 2015 ACM. 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. This material has been approved for public release and unlimited distribution. DM-0002944




Figure 1: Overview of our approach

lem, which can then be formulated as maximal specification inference. For instance, Beyene et al. [14] and Jobstmann et al. identify a set of “*suspect statements*” in a faulty program. To find a replacement for such statements, they, essentially, abstract them as unknown procedures and synthesize specifications of those procedures that ensure program correctness.

Motivated by such a diverse range of applications, we study the maximal specification inference problem in a general logic-based setting. Specifically, we propose a technique for automatically inferring maximal specifications by reducing the problem to a novel generalization of logical abduction. Given a program P containing calls to procedures $F_1, \dots, F_n$, we wish to infer a logical specification φ_i for each F_i such that the following conditions are satisfied:

- **Safety:** If each procedure F_i satisfies its specification φ_i , then program P is guaranteed to satisfy its desired safety property.
- **Non-triviality:** Our technique should not infer the trivial specification *false* for any procedure.
- **Maximality:** The specification φ_i inferred for each F_i should be a weakest such specification. That is, we cannot weaken any φ_i without either violating safety or making the specification φ_j for some other procedure F_j stronger.

In our setting, the maximality criterion (Pareto optimality) is very important because we would like to infer the weakest specifications that are needed for guaranteeing safety. Specifications that are stronger than necessary can cause a verifier to miss real bugs or a synthesizer to fail where it should not. For instance, synthesizing an overly restrictive specification of a library routine and then reusing it to verify a client application might cause us to miss bugs.

The key ingredient of our specification inference algorithm is a *counterexample-guided inductive synthesis* (CEGIS) loop with a new logical inference technique called *multi-abduction* at its core (Figure 1). Multi-abduction is a powerful generalization of *standard abduction*: Given a pair of formulas χ, C and an unknown predicate $R(\mathbf{x})$ (called an *abducible*), standard abduction infers an interpretation φ of R such that $\varphi \wedge \chi \not\models \text{false}$ and $\varphi \wedge \chi \models C$. Intuitively, standard abduction looks for an *explanatory hypothesis* φ that, together with the *known facts* χ , is sufficient to imply the *conclusion* C .

Multi-abduction generalizes standard abduction by simultaneously inferring multiple unknowns. Specifically, given a formula

$$\left(\bigwedge_{i=1}^n \bigwedge_{j=1}^{n_i} R_i(\mathbf{x}_{ij}) \wedge \chi \right) \Rightarrow C,$$

multi-abduction computes a map from each unknown predicate R_i to a formula φ_i over variables $\mathbf{v}_i$ such that

- (i) $\bigwedge_{ij} \varphi_i[\mathbf{x}_{ij}/\mathbf{v}_i] \wedge \chi \not\models \text{false}$, and
- (ii) $\bigwedge_{ij} \varphi_i[\mathbf{x}_{ij}/\mathbf{v}_i] \wedge \chi \models C$,

```

1 l ← 0
2 while (*)
3   if (p = 1)
4     l ← lock(l)
5   if (p != 1)
6     l ← 0
7   else
8     assert(l = 1)

```

Figure 2: Example program used to illustrate our technique

where $\varphi[\mathbf{x}/\mathbf{v}]$ is φ with all occurrences of $\mathbf{v}$ replaced by $\mathbf{x}$. Observe that multi-abduction generalizes standard abduction in two important ways:

- (i) The input formula is allowed to contain multiple different unknown predicates $R_1, \dots, R_n$, and
- (ii) the same predicate R_i can appear multiple times with different arguments $\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}$.

The problem of inferring unknown method specifications in a bounded program fragment is naturally encoded as multi-abduction: Each unknown specification corresponds to an unknown predicate R_i , and the same R_i can appear multiple times with different arguments at distinct call sites. Further, since we want to infer weakest specifications, we are interested in *maximal solutions* to multi-abduction problems.

As outlined in Figure 1, our approach is based on a CEGIS loop that uses multi-abduction to synthesize specifications from a finite set of bounded control-flow paths through a program P . Given a set of candidate specifications, we then invoke a verifier to check if the specifications ensure correctness of *all* of P ’s execution paths. If not, new counterexample paths are emitted by the verifier and synthesis is restarted, taking the new program paths into account.

We implemented our approach in an LLVM-based program verifier [6] and used it to synthesize maximal specifications of unknown procedures in a variety of programs. Most interestingly, we show how our approach can efficiently synthesize stubs for kernel routines and initialization code from Windows device drivers used in the software verification competition (SV-COMP) [16].

This paper makes the following key contributions:

- We define the *maximal specification inference* problem and present a *counterexample-guided inductive synthesis* algorithm for solving it.
- We define the *multi-abduction* problem, a generalization of traditional logical abduction, and present an algorithm for finding maximal solutions to multi-abduction problems.
- We show how the specification inference problem for a bounded program fragment can be formulated as multi-abduction.
- We present an implementation of our proposed techniques and demonstrate its applicability by synthesizing maximal specifications in a range of programs, including synthesizing stubs of kernel routines invoked by Windows device drivers.

2. Illustrative Example

In this section, we illustrate our maximal specification inference algorithm by applying it to a simple example.

Consider the program in Figure 2, which is a simplified version of a small locks benchmark in the software verification competition (SV-COMP) [16]. Variable l takes values in $\{0, 1\}$; the procedure `lock`, called at line 4, is unknown; the symbol $*$ represents non-deterministic choice; and the unassigned variable p holds an arbitrary value. Our goal is to synthesize a maximal specification of `lock` such that no program execution violates the assertion at line 8. Since `lock` takes one input argument and returns one value,

we are looking for a specification $lock(l_{in}, l_{out})$ that is a relation over the argument and return of `lock`, denoted l_{in} and l_{out} , respectively.

First iteration Our technique begins by setting $lock(l_{in}, l_{out})$ to $true$, the weakest possible specification. It then invokes the verifier to check if this specification is sufficient to prove the assertion. In this case, $true$ is too weak, and the verifier may return the counterexample path $\pi_1 = 1, 2, 3, 4, 5, 8$. This is a counterexample because the specification $true$ implies that `lock` at line 3 may return any value, including 0, which violates the assertion.

Now, our goal is to synthesize a specification of `lock` that ensures that π_1 is correct. To do so, we use π_1 to construct the following multi-abduction query (which, in this case, is just a classical abduction query):

$$(l_{st} = 0 \wedge p = 1 \wedge lock(l_{st}, l_{end})) \Rightarrow l_{end} = 1,$$

The left-hand side of the implication encodes the semantics of the path π_1 , where l_{st} denotes the initial (start) value of `l`, l_{end} denotes its value at the assertion, and the call to `lock` is encoded using the *abducible* $lock(l_{st}, l_{end})$. The right-hand side of this implication encodes the assertion (postcondition). Thus, this abduction query is looking for an interpretation of the abducible $lock$ that ensures the validity of this implication, and therefore the correctness of path π_1 with respect to the assertion. In this case, we compute the following maximal solution for $lock(l_{in}, l_{out}) \equiv l_{in} = 0 \Rightarrow l_{out} = 1$. This formula specifies that `lock` must return 1 whenever its input is 0; otherwise, `lock` can return an arbitrary value.

Second iteration In the second iteration, we invoke the verifier again, this time checking whether $l_{in} = 0 \Rightarrow l_{out} = 1$ is sufficient to prove correctness of the program. The verifier discovers that this specification is still too weak. As a result, it might return the counterexample $\pi_2 = 1, \langle 2, 3, 4, 5, 8 \rangle^2$, where the superscript indicates that we take the same path through the loop twice. This is a counterexample because, in the second iteration of the loop, the value of `l` is 1 and our specification allows an arbitrary return value for `lock` when the input is 1, thus violating the assertion.

We now create the following new multi-abduction problem:

$$\text{ENCODE}(\{\pi_1, \pi_2\}) \Rightarrow l_{end} = 1,$$

Here, the left-hand side of the implication encodes executions of *both* paths seen so far (i.e., π_1 and π_2), and the right-hand side is the postcondition. Specifically, $\text{ENCODE}(\{\pi_1, \pi_2\})$ corresponds to the following formula:

$$l_{st} = 0 \wedge p = 1 \wedge lock(l_{st}, l_{fst}) \wedge lock(l_{fst}, l_{snd}) \wedge l_{end} \in \{l_{fst}, l_{snd}\}.$$

Observe that the left-hand side of this multi-abduction problem has two occurrences of the unknown predicate $lock$ over different arguments, one for each iteration of the loop (fst and snd). A solution to this multi-abduction problem is an interpretation of the predicate $lock$ that makes the implication valid. Therefore, by finding a weakest solution to this multi-abduction problem, we find a maximal specification for `lock` that ensures correctness of the two paths we have examined thus far. In this case, we get the solution $lock(l_{in}, l_{out}) \equiv l_{out} = 1$.

Final iteration Finally, we invoke the verifier with the new specification $l_{out} = 1$, which is sufficient to prove correctness of our program. We thus conclude that a maximal specification for `lock` is one that always returns the value 1, no matter what the input is.

Summary Our approach synthesizes maximal specifications for loop-free program paths by solving multi-abduction problems. Furthermore, even though each iteration of the algorithm considers a finite number of control-flow paths, we can still guarantee the maximality and safety of the inferred specification for the whole program. In the following sections, we formalize this process in detail,

present our solutions, and discuss the various challenges and intricacies of developing multi-abduction solvers.

3. Maximal Specifications and Multi-Abduction

We now define the *maximal specification inference* problem and introduce *multi-abduction*. Both in this section and the rest of the paper, we assume that formulas belong to a first-order theory $\mathcal{T}$.

3.1 Maximal safe specifications

Programs A program P is a tuple $(N, E, n_e, n_x, \mathcal{V}, \rho)$, where (N, E) is a directed graph with a special *entry node* n_e and *exit node* n_x . The set $\mathcal{V}$ denotes all program variables, and ρ is a first-order formula that serves as a postcondition for P , denoted $\text{post}(P)$. We assume that each edge $e \in E$ is labelled with a command e^c , which can be one of the following constructs:

- An assignment $v \leftarrow \text{exp}$, where $v \in \mathcal{V}$ and exp is a program expression over variables $\mathcal{V}$. We assume that there is at most one assignment to each variable (i.e., the program is in SSA form).¹
- An assumption of the form $\text{assume}(b)$, where b is a Boolean expression over program variables $\mathcal{V}$.
- A procedure call $x \leftarrow F(y)$, where x and y are vectors of variables in $\mathcal{V}$ and F is a procedure. We model side effects by allowing procedures to return multiple values.

A procedure F is *unknown* if its source code is not available; we write $\text{unknowns}(P)$ to denote unknown procedures called by P . For the purposes of this paper, we will assume that all callees are unknown, and our goal is to infer safe and permissive specifications for each $F \in \text{unknowns}(P)$. We say that a program P is *closed* if it does not contain calls to unknown procedures; otherwise, it is *open*. Given a closed program P with corresponding postcondition ρ , we write $P \models \rho$ if P satisfies its postcondition in all executions. We also say that program P is *safe* whenever $P \models \rho$.

Maximal safe specifications Given program P with unknown procedures $F_1, \dots, F_n$, we use the symbol Δ to denote a *specification environment* (or *specification* for short) mapping each procedure F_i to a first-order formula φ_i over vectors of variables α, β such that α denotes F_i 's arguments and β denotes its output. We write $P[\Delta]$ to denote the closed program with every procedure call $x \leftarrow F_i(y)$ replaced by $\text{assume}(\varphi_i[x/\alpha, y/\beta])$, where $\Delta(F_i) = \varphi_i$. In other words, $P[\Delta]$ is like P but with every unknown procedure F_i behaving according to its specification $\Delta(F_i)$.

DEFINITION 1. (Safe specification) Given a program P , with $\text{post}(P) = \rho$, and a specification environment Δ , we say that Δ is a safe specification for P if $P[\Delta] \models \rho$.

While safe specifications are sufficient for ensuring safety, they may be stronger than necessary. Since we want to synthesize *weakest* safe specifications, we define a partial order on specification environments. We use $\text{dom}(\Delta)$ and $\text{range}(\Delta)$ to denote the domain and range of Δ .

DEFINITION 2. (Partial order $\succ$) We write $\Delta \succ \Delta'$, and say that Δ is weaker than Δ' , if the following hold:

- (i) $\text{dom}(\Delta) = \text{dom}(\Delta')$,
- (ii) $\forall F \in \text{dom}(\Delta). \Delta(F) \Leftarrow \Delta'(F)$, and
- (iii) $\exists F \in \text{dom}(\Delta). \Delta(F) \not\Leftarrow \Delta'(F)$.

We define *maximal safe specifications* using this partial order:

¹ For clarity of presentation, and wlog, we do not include SSA *phi* nodes in the language, as they can be encoded with *assumes* and additional variables.

DEFINITION 3. (Maximal safe specification) Given program P with $\text{post}(P) = \rho$ and a specification environment Δ^* , we say that Δ^* is a maximal safe specification for P if

- (i) $P[\Delta^*] \models \rho$ and
- (ii) for all Δ such that $\Delta \succ \Delta^*$, we have $P[\Delta] \not\models \rho$.

In other words, Δ^* is a maximal safe specification for P if we cannot weaken any $\Delta^*(F)$ while preserving the safety of $P[\Delta^*]$. Note that there may not be a unique maximal safe specification that is higher than all other safe specifications in $\succ$; instead, there may be a space of incomparable maximal safe specifications.

EXAMPLE 1. Consider the following program with postcondition $x \geq 0 \wedge z > 10$:

```
x ← F()
y ← G(x)
z ← y + 1
```

A maximal safe specification for this program is

$$\Delta : [F(\beta) \mapsto \beta \geq 0, G(\alpha, \beta) \mapsto (\alpha \geq 0 \Rightarrow \beta > 9)].$$

$\Delta(F)$ specifies that F should always return a positive value. $\Delta(G)$ specifies that G should return a value > 9 whenever its argument is positive; otherwise, it can return an arbitrary value.

3.2 Multi-abduction

Standard abduction We first review standard abduction, which we henceforth refer to as *simple abduction*. Intuitively, simple abduction asks the question, “What facts do we need to know, in addition to our already-known facts χ , in order to reach conclusion C ?” The following definition formalizes this intuition.

DEFINITION 4. (Simple abduction) Given a formula $R(\mathbf{x}) \wedge \chi \Rightarrow C$, simple abduction finds a formula φ over variables $\mathbf{x}$ such that

- (i) $\varphi \wedge \chi \not\models \text{false}$, and
- (ii) $\varphi \wedge \chi \models C$.

In this definition, we refer to $R(\mathbf{x})$ as an *abducible*. A *solution* φ to a simple abduction problem is an interpretation of $R(\mathbf{x})$ that strengthens the left-hand side of the implication in order to make the implication logically valid. A *maximal* solution to the simple abduction problem is one that is logically weakest. Every solvable simple abduction problem has a unique maximal solution up to logical equivalence.

Given a simple abduction problem $R(\mathbf{x}) \wedge \chi \Rightarrow C$, we will assume a procedure $\text{ABDUCE}(\chi, C, \mathbf{x})$ that computes a maximal solution to the corresponding simple abduction problem. For first-order theories that admit quantifier elimination, one possible implementation of ABDUCE is simply $\text{QE}(\forall \bar{\mathbf{x}}. \chi \Rightarrow C)$, where QE represents a quantifier elimination procedure and $\bar{\mathbf{x}}$ denotes all free variables in $\chi \Rightarrow C$ except $\mathbf{x}$ [30].

Multi-abduction While simple abduction allows the inference of a single unknown predicate (abducible), it is often necessary to consider multiple abducibles over different vocabularies. For example, in the context of specification inference, different unknown predicates correspond to different procedures, and different occurrences of the same unknown predicate correspond to different calls to the same procedure. Multi-abduction generalizes simple abduction by allowing multiple unknowns as well as different occurrences of the same unknown:

DEFINITION 5. (Multi-abduction) Given a formula

$$\left(\bigwedge_{i=1}^n \bigwedge_{j=1}^{n_i} R_i(\mathbf{x}_{ij}) \wedge \chi \right) \Rightarrow C,$$

```
1 function MAXSAFESPEC(P)
2    $\Delta \leftarrow \{F_i \mapsto \text{true} \mid F_i \in \text{unknowns}(P)\}$ 
3    $\Pi \leftarrow \emptyset; \hat{\Pi} \leftarrow \emptyset$ 
4   while true do
5      $\theta \leftarrow \text{VERIFY}(P[\Delta], \hat{\Pi})$ 
6     if  $\theta = \emptyset$  then return  $\Delta$ 
7      $\Pi \leftarrow \Pi \cup \theta$ 
8      $\psi \leftarrow \text{ENCODE}(\Pi)$ 
9      $\Phi \leftarrow \text{FLATTEN}(\psi)$ 
10     $\Delta' \leftarrow \text{MULTIABDUCE}(\Phi, \text{post}(P))$ 
11    if  $\Delta' = \text{none}$  then  $\hat{\Pi} \leftarrow \hat{\Pi} \cup \theta$ 
12    else  $\Delta \leftarrow \Delta'$ 
```

Algorithm 1: Synthesizing maximal safe specifications

multi-abduction finds a mapping Δ from each R_i to a formula φ_i over variables $\mathbf{x}_i$ such that

- (i) $\bigwedge_{ij} \varphi_i[\mathbf{x}_{ij}/\mathbf{x}_i] \wedge \chi \not\models \text{false}$, and
- (ii) $\bigwedge_{ij} \varphi_i[\mathbf{x}_{ij}/\mathbf{x}_i] \wedge \chi \models C$.

In this definition, we call each $R_i(\mathbf{x}_{ij})$ an abducible and refer to Δ as a solution to the multi-abduction problem. (For simplicity, we overload notation and use Δ to denote multi-abduction solutions as well as safe specifications.) A *maximal solution* to the multi-abduction problem is one where no φ_i can be weakened. In other words, Δ is a maximal solution to a multi-abduction problem $\mathcal{M}$ if, for any φ'_i that is logically weaker than $\Delta(R_i)$, the mapping $\Delta[R_i \mapsto \varphi'_i]$ is not a solution to $\mathcal{M}$. Unlike simple abduction problems, multi-abduction problems do not have unique maximal solutions. This is illustrated by the following example:

EXAMPLE 2. Consider the multi-abduction problem

$$R_1(x) \wedge R_2(y) \Rightarrow x + y \geq 1$$

and two of its solutions

$$\begin{aligned} \Delta_1 : & [R_1 \mapsto x \geq 0, R_2 \mapsto y \geq 1] \\ \Delta_2 : & [R_1 \mapsto x \geq 1, R_2 \mapsto y \geq 0]. \end{aligned}$$

In this case, both Δ_1 and Δ_2 are maximal and incomparable. $\square$

A special kind of multi-abduction problem is one where all abducibles have unique predicates—that is, each unknown R_i appears exactly once in the formula. Such problems, which we call *linear multi-abduction problems*, have the form $\bigwedge_i R_i(\mathbf{x}_i) \wedge \chi \Rightarrow C$. In contrast, we refer to the general case, where the same predicate can appear multiple times with different arguments, as *non-linear multi-abduction problems*.

4. Synthesizing Maximal Safe Specifications

In this section, we present our counterexample-guided inductive synthesis algorithm for learning maximal safe specifications.

High-level description The high-level skeleton of our technique is shown in Algorithm 1; it takes a program P and returns a maximal safe specification. The algorithm initially assumes that each $F \in \text{unknowns}(P)$ has specification *true* (line 2). While such a specification Δ is very permissive, it is unlikely to be safe. Hence, the algorithm goes through a refinement loop (lines 4–12) where each specification in Δ is iteratively modified until $P[\Delta]$ is proven safe.

Going into more detail, line 5 of Algorithm 1 attempts to verify program P (with respect to its postcondition $\text{post}(P)$) using current specifications Δ . For this purpose, we assume a procedure VERIFY that is capable of generating one or more counterexamples (as program paths) if the verification task fails. VERIFY also takes a set of paths $\hat{\Pi}$ and ensures that the returned counterexamples are

not in $\hat{\Pi}$, i.e., $\theta \cap \hat{\Pi} = \emptyset$. Since many existing verification algorithms are capable of generating counterexamples [6, 10, 17, 44], we do not describe the VERIFY procedure in detail.

Now, suppose the verification task at line 5 fails with a set θ of counterexamples that falsify $\text{post}(P)$. In this case, we add θ to the set Π of all counterexamples encountered so far and try to infer a specification environment Δ that is sufficient to rule out all counterexamples in Π .² To infer such a specification Δ , one possibility is to solve a *set* of simultaneous multi-abduction problems. That is, for each counterexample path $\pi \in \Pi$, we generate a multi-abduction problem where the abducibles correspond to procedure calls in π , and we then solve all of these multi-abduction problems *simultaneously*. Unfortunately, there are two main problems with this approach: First, it requires us to devise a technique for solving not just one multi-abduction problem, but a set of them simultaneously. Second, if we represent the set of counterexamples Π symbolically as an SMT formula ψ , generating one multi-abduction problem *per counterexample* would require us to convert ψ to DNF, a possibly exponential operation.

Encoding counterexamples In order to side-step explicit counterexample enumeration and simultaneous multi-abduction, we present a technique that encodes a set of counterexamples into a single multi-abduction problem. Specifically, we think of Π as a loop-free program P_Π and use a procedure called ENCODE to represent all executions of P_Π as a formula ψ (a standard, BMC-like encoding of loop-free programs [5, 11, 24]). Given loop-free program P_Π with nodes N and edges E , ENCODE generates the following formula:

$$c_{n_e} \wedge \bigwedge_{u \in N} \left(c_u \Rightarrow \bigvee_{e=(u,v) \in E} (c_v \wedge \llbracket e^c \rrbracket) \right)$$

Here, c_u is an auxiliary Boolean variable indicating that control is at program location u . The formula c_{n_e} indicates that the program begins execution at the entry node n_e . Now, if we reach node u , we must also reach one of its successors v and execute the command e^c associated with edge $e = (u, v)$. Hence, the formula $c_v \wedge \llbracket e^c \rrbracket$ denotes the transition from u to v , where $\llbracket e^c \rrbracket$ is defined as follows:

$$\llbracket x \leftarrow e \rrbracket = (x = e) \quad \llbracket \text{assume}(b) \rrbracket = b$$

$$\llbracket x \leftarrow F_i(\mathbf{y}) \rrbracket = R_i(\mathbf{x}, \mathbf{y})$$

Our encoding of assignments assumes that the program has been converted to SSA form, as we do not explicitly rename variables. Also, since we assume that all callees F_i are unknown, we introduce a predicate R_i that represents the unknown specification of F_i .

Going back to Algorithm 1, let ψ be a formula encoding all counterexamples Π encountered so far. To ensure that each $\pi \in \Pi$ is ruled out in the next iteration, we must find an interpretation of all predicates R_i under which the formula $\psi \Rightarrow \text{post}(P)$ is logically valid. Observe that this is almost a multi-abduction problem, but not quite, as the abducibles R_i can appear under disjunctions in ψ , whereas multi-abduction requires formula ψ to be of the form $\bigwedge_{i,j} R_i(\mathbf{x}_{ij}) \wedge \chi$ (i.e., all abducibles appear at the outermost conjunction). Fortunately, given a formula ψ where all unknowns appear positively (i.e., under an even number of negations), we can always transform ψ to the syntactic form $\bigwedge_{i,j} R_i(\mathbf{x}_{ij}) \wedge \chi$ using a procedure we call *flattening*.

Multi-abduction problem flattening Given an arbitrary formula φ where unknown predicates appear only positively, the procedure FLATTEN, given in Algorithm 2, shows how to convert φ to a multi-abduction problem. Specifically, procedure FLATTEN takes as input a formula φ , where all abducibles appear only positively in φ , and

```

1 function FLATTEN( $\varphi$ )
2    $\varphi^* \leftarrow \varphi; \psi \leftarrow \text{true}$ 
3   for all  $R(\mathbf{x}) \in \text{abds}(\varphi^*)$  do
4      $\mathbf{x}_i \leftarrow \text{fresh}(\mathbf{x})$ 
5      $\varphi^* \leftarrow \varphi^*[(\mathbf{x} = \mathbf{x}_i)/R(\mathbf{x})]$ 
6      $\psi \leftarrow \psi \wedge R(\mathbf{x}_i)$ 
7   return  $\psi \wedge \varphi^*$ 

```

Algorithm 2: Extracting abducibles to the outer level

returns a formula of the form $\bigwedge_{i,j} R_i(\mathbf{x}_{ij}) \wedge \chi$, where all abducibles have been *pulled out*.

The main idea behind FLATTEN is to replace every abducible $R(\mathbf{x})$ in φ with the formula $\mathbf{x} = \mathbf{x}_i$ where $\mathbf{x}_i$ is a fresh set of variables. After this transformation has been applied to φ for every $R_i(\mathbf{x}_{ij})$, we obtain a formula φ^* that does not contain any unknown predicates (line 5). To obtain the final *flattened* formula, we then generate

$$\varphi^* \wedge \bigwedge_{R(\mathbf{x}) \in \text{abd}(\varphi)} R(\mathbf{x}_i),$$

where $\mathbf{x}_i$ is the fresh set of variables associated with abducible $R(\mathbf{x})$ and $\text{abd}(\varphi)$ is the set of all abducibles appearing in φ .

EXAMPLE 3. Consider the formula $\varphi : G(a, b) \vee a = 3 \vee G(c, d)$. Applying FLATTEN to φ yields the formula: $\psi \wedge \varphi^*$, where

$$\begin{aligned} \psi &: G(x_1, x_2) \wedge G(x_3, x_4) \\ \varphi^* &: (x_1 = a \wedge x_2 = b) \vee a = 3 \vee (x_3 = c \wedge x_4 = d) \quad \square \end{aligned}$$

The following theorem states the correctness of FLATTEN by showing that it preserves non-trivial, maximal, safe solutions, where a non-trivial solution does not set any abducible to *false*.

THEOREM 1 (Correctness of FLATTEN). *Let φ be a formula containing abducibles appearing only positively. Let C be an abducible-free formula. Then,*

$$\varphi[\Delta] \models C \text{ iff } \text{FLATTEN}(\varphi)[\Delta] \models C,$$

where Δ is a non-trivial assignment of abducibles to formulas.

PROOF 1. All proofs are available in Appendix A.

Correctness Going back to Algorithm 1, the formula Φ obtained at line 9 encodes all counterexamples Π in a form where all unknowns have been pulled out and has the syntactic form $\bigwedge_{i,j} R_i(\mathbf{x}_{ij}) \wedge \chi$. Since $\Phi \Rightarrow \text{post}(P)$ now corresponds to a multi-abduction problem, we can compute a maximal solution for the unknown predicates of Φ using the algorithm described in Section 5.

Now, if the call to MULTIABDUCE at line 10 returns a solution, we obtain a new specification environment Δ and continue our refinement loop. However, if the multi-abduction problem does not have a solution (i.e., $\Delta' = \text{none}$ at line 11), this means that the multi-abduction problem only has solutions that make every path in Π infeasible. Hence, we add the most recent counterexamples θ to set $\hat{\Pi}$; this just ensures that VERIFY does not return a path in $\hat{\Pi}$ in the next iteration. $\hat{\Pi}$ is maintained simply to ensure that we do not synthesize a specification that makes all program paths infeasible. Appendix C illustrates the need for $\hat{\Pi}$ with an example.

The following theorem states correctness of MAXSAFESPEC. Given the undecidability of the problem, it is not guaranteed to terminate in general. However, for programs with variables over finite domains (e.g., Boolean programs), it will eventually terminate.

THEOREM 2 (Correctness of MAXSAFESPEC). *If MAXSAFESPEC terminates on program P , the computed specification environment Δ is a safe, maximal, non-trivial specification for P .*

²The careful reader may wonder why we need to consider all paths in Π in every iteration. For an answer, we direct the reader to Appendix B

```

1 function MULTIABDUCELINEAR( $\chi, C, \mathcal{A}$ )
2    $\psi \leftarrow \text{ABDUCE}(\chi, C, \text{vars}(\mathcal{A}))$ 
3   if  $\psi = \text{none}$  then return none
4   return CARTDECOMP( $\chi, \psi, \mathcal{A}$ )

5 function CARTDECOMP( $\chi, \psi, \mathcal{A}$ )
6    $\Delta \leftarrow \text{INITSOLN}(\chi, \psi, \mathcal{A})$ 
7   for all  $R_i(\mathbf{x}_i) \in \mathcal{A}$  do
8      $\Delta(R_i) \leftarrow \text{ABDUCE}(\bigwedge_{j \neq i} \Delta(R_j), \psi, \mathbf{x}_i)$ 
9   return  $\Delta$ 

10 function INITSOLN( $\chi, \psi, \mathcal{A}$ )
11    $M \leftarrow \text{model}(\chi \wedge \psi)$ 
12   for all  $R_i(\mathbf{x}_i) \in \mathcal{A}$  do
13      $\Delta(R_i) \leftarrow (\mathbf{x}_i = M(\mathbf{x}_i))$ 
14   return  $\Delta$ 

```

Algorithm 3: Solving linear multi-abduction problems

5. Multi-abduction Algorithm

In this section, we present an algorithm for computing maximal solutions to multi-abduction problems.

5.1 Solving linear multi-abduction problems

Before we present our multi-abduction algorithm in its full generality, we will first present an algorithm for solving linear multi-abduction problems. Recall from Section 3 that a linear multi-abduction problem is of the form $\bigwedge_i R_i(\mathbf{x}_i) \wedge \chi \Rightarrow C$ (i.e., each unknown appears only once in the formula). While the full algorithm is a natural generalization of this one, the algorithm we present in this section is useful for gaining intuition.

Algorithm 3 summarizes our approach for solving linear multi-abduction problems. The procedure MULTIABDUCELINEAR takes as input formulas χ and C , which correspond to the unknown-free formulas on the left- and right-hand sides of the implication in Definition 5, respectively. MULTIABDUCELINEAR also takes as input a set $\mathcal{A}$ of abducibles of the form $R_i(\mathbf{x}_i)$ and returns a maximal solution Δ to the problem. We use $\text{vars}(\mathcal{A})$ to denote the set of all arguments of all abducibles in $\mathcal{A}$. The high-level structure of our multi-abduction algorithm consists of two steps:

1. **Simple abduction:** First, we solve the simple abduction problem defined by $R^*(\mathbf{x}) \wedge \chi \Rightarrow C$, where $\mathbf{x}$ denotes all variables appearing in the abducibles. In other words, we combine all abducibles into one abducible R^* and find a monolithic formula that, informally speaking, *contains* all solutions of the multi-abduction problem. We assume that the ABDUCE procedure returns none if the simple abduction problem has no solution.
2. **Cartesian decomposition:** Given a solution ψ to the simple abduction problem from step (1), we decompose ψ into a set of formulas $\varphi_1, \dots, \varphi_n$ such that
 - (i) each φ_i is only over $\mathbf{x}_i$,
 - (ii) $\bigwedge_i \varphi_i \Rightarrow \psi$, and
 - (iii) the set of formulas $\varphi_1, \dots, \varphi_n$ is as weak as possible.

We refer to the decomposition of formula ψ into such a set $\varphi_1, \dots, \varphi_n$ as *Cartesian decomposition*—intuitively, we are decomposing ψ into a Cartesian product of the formulas φ_i .

Since we already know how to solve simple abduction problems (Section 3.2), the key part of our multi-abduction algorithm is the Cartesian decomposition function (CARTDECOMP) of Algorithm 3 (line 5). This procedure first finds an initial (non-maximal) solution $\Delta : [R_i \mapsto \varphi_i]$ (line 6) and then iteratively weakens this solution (lines 7-8) until we obtain a maximal one.

The task of finding an initial solution Δ is performed by the INITSOLN procedure. Specifically, INITSOLN finds an interpretation of each R_i of the form $\mathbf{x}_i = \mathbf{m}_i$, where $\mathbf{m}_i$ is a vector of constants. Hence, our initial solution corresponds to a concrete assignment or *point*. To obtain such a solution, we simply get a model M of the formula $\chi \wedge \psi$ and assign each R_i to $\mathbf{x}_i = M(\mathbf{x}_i)$, where $M(\mathbf{x}_i)$ is the vector of assignments to variables $\mathbf{x}_i$ in the model M . Since this interpretation is consistent with χ , and since $\bigwedge_i \Delta(R_i) \Rightarrow \psi$, we know that Δ is a solution to the multi-abduction problem, albeit not necessarily a maximal one.

Now, the loop in lines 7-8 of the CARTDECOMP function iteratively weakens the solution returned by INITSOLN, with the goal of making it maximal. To weaken the solution for each R_i , we *fix* the solution for all the other R_j 's and weaken R_i as much as possible by solving the following simple abduction problem:

$$(R_i(\mathbf{x}_i) \wedge \bigwedge_{j \neq i} \Delta(R_j)) \Rightarrow \psi.$$

In other words, we use the existing solution given by Δ for all other R_j 's and infer the weakest R_i that still implies ψ . This strategy ensures that Δ remains a valid solution to our multi-abduction problem in every iteration.

EXAMPLE 4. Consider the following multi-abduction problem:

$$R_1(x) \wedge R_2(y) \wedge x \geq 0 \Rightarrow (x + y \geq 1 \vee z \geq 0).$$

The algorithm starts by solving the following simple abduction problem:

$$R^*(x, y) \wedge x \geq 0 \Rightarrow (x + y \geq 1 \vee z \geq 0).$$

This problem has maximal solution $R^* \mapsto (x \geq 0 \Rightarrow x + y \geq 1)$. Next, using INITSOLN, we get a model of this formula that is consistent with $x \geq 0$, say $x = 0$ and $y = 1$. Hence, the initial solution is $R_1 \mapsto x = 0, R_2 \mapsto y = 1$. Now, we fix R_2 to $y = 1$ and try to weaken R_1 by solving the simple abduction problem:

$$(R_1(x) \wedge y = 1 \wedge x \geq 0) \Rightarrow x + y \geq 1.$$

The answer to this problem is $R_1 \mapsto \text{true}$; hence, our current solution becomes $R_1 \mapsto \text{true}, R_2 \mapsto y = 1$. Now, we try to weaken R_2 by fixing R_1 and solving the simple abduction problem:

$$(\text{true} \wedge R_2(y) \wedge x \geq 0) \Rightarrow x + y \geq 1.$$

Since ABDUCE yields $y \geq 1$, the final solution is

$$R_1 \mapsto \text{true}, R_2 \mapsto y \geq 1.$$

The following theorem states soundness and completeness of our multi-abduction procedure.

THEOREM 3 (Correctness of MULTIABDUCELINEAR). *If the multi-abduction problem is solvable, Algorithm 3 terminates with a maximal solution; otherwise, it returns none.*

5.2 Solving non-linear multi-abduction problems

We now consider our algorithm for solving non-linear multi-abduction problems (see Algorithm 4). The general structure of this algorithm is similar to Algorithm 3 in that we first solve a simple abduction problem and then apply Cartesian decomposition. However, since we must assign the *same interpretation* to two different abducibles $R_i(\mathbf{x}_{ij})$ and $R_i(\mathbf{x}_{ik})$, the CARTDECOMP function of Algorithm 4 is somewhat more involved.

Like in the linear case, the main procedure MULTIABDUCE takes as input formulas χ and C and a set of abducibles $\mathcal{A}$ of the form $R_i(\mathbf{x}_{ij})$ —i.e., $\mathcal{A}$ potentially contains several occurrences of the same unknown R_i . We use $\text{preds}(\mathcal{A})$ to denote the set of unique predicates $R_1, \dots, R_n$ appearing in $\mathcal{A}$.

Cartesian decomposition As in Algorithm 3, the CARTDECOMP procedure first computes an initial non-maximal solution Δ for each abducible R_i (line 6) and then iteratively weakens Δ to obtain a maximal solution (loop at line 7). The difference here is in the way CARTDECOMP treats repeated occurrences of the same unknown. The weakening loop at line 7 works in two steps:

1. In line 8, we group all occurrences of each R_i as one monolithic abducible R_i^* over the variables $\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}$, and computes a maximal solution for R_i^* . This is done using the simple abduction query:

$$\left(R_i^*(\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}) \wedge \bigwedge_{k \neq i} \bigwedge_j \Delta(R_k)[\mathbf{x}_{kj}/\mathbf{x}_j] \right) \Rightarrow \psi.$$

The computed solution ψ_i at line 8 is the weakest formula over variables $\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i}$ that makes the above implication valid.

2. At this point, we need to decompose ψ_i into n_i formulas, one for each occurrence of R_i . These formulas also have to be equivalent, as per definition of the multi-abduction problem. More formally, we would like to find a weakest formula φ_i over the fresh vector of variables $\mathbf{x}_i$ such that: $\bigwedge_j \varphi_i[\mathbf{x}_{ij}/\mathbf{x}_i] \Rightarrow \psi_i$. We call this process *isomorphic decomposition* (line 9), as it decomposes the formula ψ_i into a conjunction of formulas $\varphi_i[\mathbf{x}_{ij}/\mathbf{x}_i]$ that are equivalent modulo variable renaming.

Initial solution Since the main difference of Algorithm 4 from MULTIABDUCELINEAR lies in isomorphic decomposition and initial solution inference, we now turn our attention to lines 11-27 of the algorithm. Let us first consider the INITSOLN procedure for finding an initial non-maximal interpretation for each unknown R_i . First of all, if the multi-abduction problem is to have a solution, then there must be some interpretation of R_i of the form $\bigvee_j \mathbf{x}_i = \mathbf{m}_{ij}$, where each $\mathbf{m}_{ij}$ is a vector of constants. Note that we do not require these $\mathbf{m}_{ij}$'s to be distinct; hence, even if the only interpretation of R_i is a single *point*, we can still find such a solution.

Now, to find a solution of this form, we need to find values of $\mathbf{m}_{ij}$ under which the formula

$$\left(\bigwedge_{ij} \bigvee_k \mathbf{x}_{ij} = \mathbf{m}_{ik} \right) \Rightarrow \psi$$

is valid. Specifically, we let $\mathbf{m}_{ij}$ be fresh symbolic constants and ask for a model M of the formula shown in line 13, where $\mathcal{A}(R_i)$ denotes all abducibles containing unknown R_i and $\text{vars}(\psi)$ is the set of free variables in ψ . Then, we obtain an initial solution for $\Delta(R_i)$ as $\bigvee_j \mathbf{x}_i = M(\mathbf{m}_{ij})$. The intuition here is as follows: In the linear case, INITSOLN initializes every abducible R_i with an assignment $\mathbf{x}_i = \mathbf{m}_i$. In the non-linear case, we do the same, but we need to make sure that if we initialize one occurrence of R_i with some assignment A , then every other occurrence should also have that assignment. This is captured by the formula in line 12.

Isomorphic decomposition Let us now consider the ISODECOMP procedure in line 17. Given a formula ψ and a set of abducibles $\mathcal{R} = \{R_i(\mathbf{x}_{i1}), \dots, R_i(\mathbf{x}_{in_i})\}$, recall that the goal of isomorphic decomposition is to find a weakest formula φ such that

$$\bigwedge_j \varphi[\mathbf{x}_{ij}/\mathbf{x}_i] \Rightarrow \psi \quad (1)$$

is valid. To find such a formula φ , ISODECOMP first initializes φ to the current solution of R_i given by Δ . Note that, by construction of the initial solution and the maximality guarantee of ABDUCE, $\Delta(R_i)$ always satisfies Formula 1. Now, starting from this initial solution at line 18, ISODECOMP iteratively weakens φ until it can no longer be weakened (lines 19-27).

```

1 function MULTIABDUCE( $\mathcal{X}, C, \mathcal{A}$ )
2    $\psi \leftarrow \text{ABDUCE}(\mathcal{X}, C, \text{vars}(\mathcal{A}))$ 
3   if  $\psi = \text{none}$  then return none
4   return CARTDECOMP( $\mathcal{X}, \psi, \mathcal{A}$ )

5 function CARTDECOMP( $\mathcal{X}, \psi, \mathcal{A}$ )
6    $\Delta \leftarrow \text{INITSOLN}(\mathcal{X}, \psi, \mathcal{A})$ 
7   for all  $R_i \in \text{preds}(\mathcal{A})$  do
8      $\psi_i \leftarrow \text{ABDUCE}(\bigwedge_{k \neq i} \bigwedge_j \Delta(R_k)[\mathbf{x}_{kj}/\mathbf{x}_j], \psi, \mathbf{x}_i)$ 
9      $\Delta(R_i) \leftarrow \text{ISODECOMP}(\psi_i, \mathcal{A}(R_i), \Delta)$ 
10  return  $\Delta$ 

11 function INITSOLN( $\mathcal{X}, \psi, \mathcal{A}$ )
12   $\psi_M \leftarrow \forall \text{vars}(\psi). \left( \bigwedge_{R_i(\mathbf{x}_{ij}) \in \mathcal{A}} \bigvee_{k=1}^{|\mathcal{A}(R_i)|} \mathbf{x}_{ij} = \mathbf{m}_{ik} \right) \Rightarrow \psi$ 
13   $M \leftarrow \text{model}(\mathcal{X}[\mathbf{m}_{ij}/\mathbf{x}_{ij}] \wedge \psi_M)$ 
14  for all  $R_i \in \text{preds}(\mathcal{A})$  do
15     $\Delta(R_i) \leftarrow \bigvee_j \mathbf{x}_i = M(\mathbf{m}_{ij})$ 
16  return  $\Delta$ 

17 function ISODECOMP( $\psi, \mathcal{R}, \Delta$ )
18   $\varphi \leftarrow \Delta(\text{preds}(\mathcal{R}))$ 
19  while true do
20     $\psi_M \leftarrow \forall \text{vars}(\psi). \left( \bigwedge_{R(\mathbf{x}_{ij}) \in \mathcal{R}} \bigvee_{k=1}^{|\mathcal{R}|} \mathbf{x}_{ij} = \mathbf{m}_{ik} \right) \Rightarrow \psi$ 
21     $M \leftarrow \text{model} \left( \bigwedge_{j=1}^{|\mathcal{R}|} \neg \varphi[\mathbf{m}_{ij}/\mathbf{x}_i] \wedge \psi_M \right)$ 
22    if  $M = \emptyset$  then return  $\varphi$ 
23     $\varphi \leftarrow \varphi \vee \bigvee_j \mathbf{x}_i = M(\mathbf{m}_{ij})$ 
24     $\varphi_j \leftarrow \varphi[\mathbf{x}_{ij}/\mathbf{x}_i]$  for each  $j$ 
25    for all  $R(\mathbf{x}_{ij}) \in \mathcal{R}$  do
26       $\varphi_j \leftarrow \text{ABDUCE}(\bigwedge_{k \neq j} \varphi_k, \psi, \mathbf{x}_{ij})$ 
27     $\varphi \leftarrow \bigwedge_j \varphi_j[\mathbf{x}_i/\mathbf{x}_{ij}]$ 

```

Algorithm 4: Solving non-linear multi-abduction problems

In each iteration of the while loop, we weaken our current solution φ by finding a *set* Θ of assignments $\mathbf{x}_i = \mathbf{m}_{ij}$ such that (i) $\mathbf{x}_i = \mathbf{m}_{ij}$ currently does not satisfy φ , and (ii) adding Θ to the satisfying assignments of φ does not violate Formula 1. If such an assignment does not exist, this means φ is already the weakest solution that makes Formula 1 valid; hence, we return φ (line 22). Otherwise, we weaken our current solution φ by adding models Θ (loop at line 25).

Lines 24-27 of the algorithm are heuristics to accelerate convergence. In particular, since the ISODECOMP algorithm (as described so far) adds a finite, and typically small, set of models in each iteration, we try to further weaken φ using simple abduction. Specifically, at line 26, we ask, “How much can we weaken $R_i(\mathbf{x}_{ij})$ while fixing the other $R_i(\mathbf{x}_{ik})$ ’s and maintaining the validity of Formula 1?” Hence, the formula φ_j obtained at line 27 has the property that $\varphi \Rightarrow \varphi_j[\mathbf{x}_i/\mathbf{x}_{ij}]$. Finally, since all abducibles $R_i(\mathbf{x}_{ij})$ must have the same interpretation modulo renaming, we obtain a new solution φ as $\bigwedge_j \varphi_j[\mathbf{x}_i/\mathbf{x}_{ij}]$ (line 27). Note that the formula φ at line 27 is at least as weak as the version of φ at line 23, although it may not be *strictly* weaker.

EXAMPLE 5. Consider the multi-abduction problem

$$R(x) \wedge R(y) \Rightarrow x + y \geq 2$$

in the first-order theory of linear integer arithmetic (LIA). The algorithm calls CARTDECOMP on the formula $x + y \geq 2$, and, regardless of what INIT SOLN returns, it calls ISODECOMP with $\psi = x + y \geq 2$ and abducibles $\mathcal{R} = \{R(x), R(y)\}$. Now, assuming INIT SOLN sets $R(\alpha) \equiv \alpha = 2 \vee \alpha = 3$, we weaken the solution φ by computing a new model: Suppose that φ at line 23 is now assigned to $\alpha = 2 \vee \alpha = 3 \vee \alpha = 4$. In the loop at line 25, we first use simple abduction to weaken $R(x)$, which yields $x \geq 0$. In the next iteration, we weaken $R(y)$ to $y \geq 2$; hence the new φ becomes $\alpha \geq 2$ at line 27. In the next iteration of the while loop, we ask for a new model of x and y and obtain $x = 1$, $y = 1$. Hence, the new φ at line 23 becomes $\alpha \geq 2 \vee \alpha = 1$. When we use simple abduction to weaken $R(x)$ in line 26, we now obtain $R(x) \equiv x \geq 1$. Similarly, when we weaken $R(y)$ using $R(x) \equiv x \geq 1$, we obtain $y \geq 1$ for $R(y)$. Hence, the final solution is $R(\alpha) \equiv \alpha \geq 1$. $\square$

Correctness and termination The following theorem states soundness of the MULTIABDUCE algorithm.

THEOREM 4 (Correctness of MULTIABDUCE). *If MULTIABDUCE returns solution Δ , then Δ is a maximal solution; if it returns none, then there is no solution to the multi-abduction problem.*

While the MULTIABDUCE procedure of Algorithm 4 always terminates for logical theories with finite ascending chains of implications (e.g., propositional logic), it does not have termination guarantees otherwise. This is expected because multi-abduction problems in some theories are not guaranteed to have maximal solutions that are finitely expressible in that theory. For example, consider the multi-abduction problem $R(x) \wedge R(y) \Rightarrow y \neq x + 1$ in linear real arithmetic. A maximal solution for $R(\alpha)$ is $\bigvee_{k \in \mathbb{Z}} (2k < \alpha \leq 2k + 1)$, which is not finitely expressible in this theory and requires integral constraints. We formalize lack of maximal solutions for non-linear multi-abduction in the following theorem.

THEOREM 5 (Maximality in FOL theories). *There are first-order theories for which there exists non-linear multi-abduction problems that have no maximal solutions.*

6. Implementation and Extensions

6.1 Implementation

We implemented a prototype of our maximal specification synthesis algorithm (MAXSAFESPEC) and its multi-abduction core as an extension of UFO [6], an open-source, LLVM-based [2] verification framework for C code. We extended UFO’s front end to enable procedure calls with multiple returns, just as in our program model from Section 3. This model allows us to handle procedures that can modify global state.

To compute inductive invariants, UFO unrolls a program’s CFG into a DAG and uses DAG *interpolants* [4] to hypothesize an inductive invariant in case all paths in the unrolling are safe. Our implementation utilizes the DAG-shaped counterexamples emitted by UFO in order to drive the specification synthesis loop. Of course, we could instrument UFO to emit a single counterexample path at a time—instead of a DAG of counterexamples. However, we found that this process does not scale at all as soon as we go beyond very simple programs, due to the large number of paths that need to be examined in order to arrive at a maximal safe specification.

Our multi-abduction algorithm implementation uses the Z3 SMT solver [28] for satisfiability checking and quantifier elimination. Programs are encoded in linear integer arithmetic (QLIA).

6.2 Optimizations

To produce an efficient implementation of our specification synthesis algorithm, we identified a number of key optimizations.

Aggressive formula preprocessing Unsurprisingly, the most expensive operation in our procedure is the universal quantifier elimination step needed for simple abduction. In our implementation context, we have to apply quantifier elimination to large formulas over linear integer arithmetic (also known as Presburger arithmetic). It is well known that the worst case space/time bound for this problem is super exponential. In addition to this theoretical intractability, quantifier elimination has not received the same attention—in terms of algorithmic and engineering advances—as did satisfiability. To make matters even worse, we are dealing with large formulas encoding CFG unrollings, whereas typical uses of quantifier elimination are restricted to relatively small formulas (e.g., encoding sets of states in some numerical abstract domain).

In our initial attempts, we implemented a simple abduction procedure that directly passes the formula to Z3’s quantifier elimination procedure [19]—a combination of Cooper’s algorithm [25] and the Omega test [46]. This turned out to be very impractical. As such, we devised a two-step preprocessing phase that (i) aggressively simplifies the formula by identifying and removing redundant subformulas and (ii) attempts to eliminate *easy* variables by examining the formula syntactically. After performing these two steps, we fall back on Z3 to eliminate the remaining variables from the simplified formula.

First, to simplify a formula φ into an equivalent formula φ_s , we exploit UNSAT cores. Specifically, we ask the SMT solver, “Which literals in φ can we replace with *true* without weakening the formula?” We encode this question as follows: First, we create a formula φ' that is a negation normal form (NNF) version of φ with every literal l_i replaced by $a_i \Rightarrow l_i$, where a_i is a fresh Boolean *assumption* variable associated with literal l_i . Note that if we set all a_i variables in φ' to *true*, we get a formula equivalent to φ . Now, we ask Z3 whether the formula

$$\varphi' \wedge \neg \varphi \wedge \bigwedge_i a_i$$

is unsatisfiable. This formula is unsatisfiable iff $\varphi \Rightarrow \varphi$ is valid, which is obviously the case. As a result, we can extract an UNSAT core over the assumption variables $\{a_i\}$; the UNSAT core indicates which literals in φ can be replaced by *true* without changing its meaning. As Z3 is not guaranteed to return a minimal UNSAT core, we invoke it multiple times to get rid of as many literals as possible.

Second, after simplification, we examine the formula syntactically to eliminate as many variables as possible before passing the formula on to Z3. For instance, consider the formula

$$\forall x. x = y + z \vee \varphi,$$

where x does not appear in φ . In this case, we know that the literal $x = y + z$ can be replaced by *false*, thus eliminating the quantified variable x . We found that this simple heuristic eliminates a substantial number of variables, simplifying quantifier elimination for Z3.

Aggressive result postprocessing In addition to preprocessing the input to the quantifier elimination procedure, we also simplify the output of quantifier elimination. Given the complexity of quantifier elimination in LIA, the resulting formulas can be quite large and, in our experience, full of redundant subformulas.

Since our goal is to compute procedure specifications, we would like to produce simple specifications that are easy to examine by a human and/or efficient for reuse in verification. To achieve that, we applied the aforementioned UNSAT core simplification technique to get rid of redundant literals. For postprocessing, we go even further with simplification: we traverse the formula and attempt to remove redundant subformulas using local SMT checks.

Finding initial solutions The key component of our multi-abduction procedure is Cartesian decomposition. The maximal solution computed by Cartesian decomposition is determined by

the initial solution (computed by INIT SOLN). To compute *higher quality* specifications, we attempt to compute initial solutions that maximize the number of feasible paths. This ensures that computed specifications are not only maximal, but also permit as many program paths as possible. To do this, we implemented a greedy MAXSMT procedure that finds a model that maximizes the number of satisfiable disjuncts (program paths) in the formula χ in Algorithms 3 and 4.

Termination and non-linear multi-abduction Recall that in the presence of repeated occurrences of the same abducible (non-linear case), MULTIABDUCE may not terminate. For our set of benchmarks, the heuristic abduction loop, line 26 of Algorithm 4, always forced our algorithm to find a maximal solution. That is, we did not encounter non-terminating multi-abduction problems in practice.

7. Evaluation

Our evaluation is designed to assess (i) the efficiency and performance of our multi-abduction procedure and (ii) the quality of the synthesized specifications. For these purposes, we performed the following two sets of experiments:

- E1** Our first experiment is primarily designed to evaluate and study the performance of our multi-abduction procedure. We apply the multi-abduction procedure to problems generated by (i) a parameterized set of microbenchmarks of growing size and (ii) a set of standard verification benchmarks that we use in a program repair scenario.
- E2** Our second experiment evaluates the performance of our algorithm and the quality of synthesized specifications. For this, we used a collection of Windows device drivers from the software verification competition, SV-COMP [16]. For these programs, we eliminated kernel routine and initialization stubs, and synthesized maximal specifications that maintain correctness of the drivers. Stubs are typically manually written to abstract OS routines and avoid analysis of the whole OS kernel in the process of analyzing a single driver. Our evaluation demonstrates the potential of our approach at automatically synthesizing specifications of kernel routines from their usage context. One potential scenario would be that human auditors can then certify such specifications for reuse in verification of other drivers.

To the best of our knowledge, no existing technique can synthesize maximal safe specifications for the kinds of programs we address here. For our benchmarks, synthesizing maximal specifications requires multi-abduction and cannot be performed using simple abduction alone. Specifically, our benchmarks (i) call the same procedure multiple times and/or (ii) call different procedures. Section 8 provides a detailed comparison with related works.

7.1 E1: Evaluating multi-abduction

We now discuss our evaluation of multi-abduction. We first present our results on parameterized microbenchmarks and then on the locks benchmarks from SV-COMP.

7.1.1 Parameterized microbenchmarks

Benchmarks description We created four microbenchmarks that are parameterized by the number of calls to unknown procedures. All four benchmarks involve linear arithmetic operations and exhibit different scenarios for our tool. We evaluated our implementation on programs of increasing size. Our goal with this set of benchmarks is to stress our multi-abduction technique and experiment with how far it can go by increasing input sizes. We first describe our microbenchmarks and the synthesized specifications.

Consider the accumulator benchmark in the top part of Figure 3. In every iteration of the loop, the unknown procedure `fun`

returns the values of x and y , which are then added to z . This benchmark is parameterized by the number of times, n , lines l_1 and l_2 are repeated within the loop. The loop executes arbitrarily many times (as denoted by the condition $*$). Our goal is to find a specification for `fun` that ensures that $z > 0$ after the loop exit. For every instantiation of this benchmark, our tool computes the maximal specification $r_1 + r_2 \geq 0$, where r_1 and r_2 denote the first and second return values of `fun`.

Unlike the accumulator benchmark, `sum` and `sum-arg` require synthesizing specifications of increasing size as we increase the parameter n . For `sum`, n determines the number of `fun`'s return values; our tool computes the maximal specification

$$\sum_{i=1}^n i * r_i \geq 10,$$

where r_i denotes the i 'th return value of `fun`. Hence, the inferred specification ensures that $y \geq 10$ throughout execution.

The benchmark `sum-arg` is trickier. Here, we call `fun` n times with n different arguments, from 1 to n , and we want to ensure that the sum of all return values is 0. One might be tempted to say that a maximal specification is $r = 0$, that is, `fun` always returns 0. This is too strong, as we only care about the return value of `fun` when its argument is between 1 and n . For each instantiation of this benchmark, our tool computes the maximal specification

$$\bigvee_{i=1}^n arg = i \Rightarrow r = 0,$$

where arg denotes the input argument of `fun` and r represents its return value. This specification states that `fun` returns 0 when its input is between 1 and n ; otherwise, it can return any value.

Finally, the benchmark `pareto` is similar to `sum`, but sets the values of $x_1, \dots, x_n$ by making calls to n different procedures. In this case, there are many maximal specifications for `fun_1, \dots, fun_n` that are incomparable as per our partial order (Definition 2). Our tool thus returns one possible maximal specification of the form

$$[\text{fun}_1 \mapsto r_1 \geq c_1, \dots, \text{fun}_n \mapsto r_n \geq c_n],$$

where r_i is the return value of `fun_i`, $c_i \in \mathbb{Z}$, and the following constraint is satisfied: $\sum_{i=1}^n i * c_i \geq 10$.

Results For each of the four benchmarks, we evaluated the performance of our multi-abduction procedure on the generated multi-abduction problem. (For these benchmarks, the multi-abduction solver is invoked once, as UFO always provides a sufficient DAG of counterexamples.) As we increase the size of each benchmark program, by increasing the parameter n , the size of the multi-abduction problem increases linearly. We measure formula size as the number of Boolean connectives in the DAG representation of the formula. Similarly, the numbers of variables and abducibles increase linearly with n .

The bottom part of Figure 3 represents the multi-abduction run-time results for each benchmark as the size of the formula increases. The time limit per benchmark is 3 minutes. Points on the top border are benchmarks that did not complete within time limit. We use two quantifier elimination strategies: (i) eliminating one variable at a time, by multiple calls to Z3, and (ii) variable partitioning—implemented internally in Z3 [19]. For each parameterized benchmark, we report results from the faster strategy.

Overall, our results indicate our multi-abduction technique's ability to scale to large formulas, with thousands of Boolean connectives, representing encodings of unrolled CFGs. For all of the formulas here, the number of quantified variables in the first call to ABDUCE (line 2 of Algorithm 4) is in the 100s. For instance, for formulas of size ~ 2000 from the accumulator benchmark, the number of abducibles is ~ 200 , and the number of quantified vari-




Figure 3: Parameterized microbenchmarks and their results. Time is in seconds.




Figure 4: Results for locks benchmarks

ables is ~ 400 (out of a total of ~ 800 variables). In `sum-arg`, for formulas of size ~ 1500 , we see ~ 200 abducibles and ~ 100 quantified variables (out of a total of ~ 500 variables). Notice the need for multi-abduction for these benchmarks: The same procedure `fun` is called multiple times, and there are calls to different procedures `funi`. Thus, we require Cartesian and isomorphic decomposition to compute specifications of unknown procedures.

7.1.2 Locks benchmarks

Benchmark description We also evaluated our approach on the locks benchmarks of SV-COMP (from the locks group), a set of benchmarks designed as challenge programs for automated software verifiers [18]. These benchmarks non-deterministically acquire and release a number of locks in a loop. The number of paths in the program increases exponentially in the number of locks. For each lock, we abstracted the code that checks whether the lock is acquired as a *unique* unknown procedure and used our tool to synthesize a specification under which the program is correct. Hence, the number of unknown procedures in these benchmarks is equal to the number of locks. The setup here mimics a *program repair* scenario: The assumption is that the code that checks whether the lock is acquired is “*suspect*” [37], and we are asking, “*What can we replace this code with in order to make the program correct?*”

Our tool synthesized a weaker specification for the abstracted code, instilling more non-determinism in these benchmarks while maintaining their correctness. Specifically, the checking code for each lock lk_i returns *true* if lk_i is not equal to 1, and *false* otherwise. Our tool returns the slightly weaker specification

$$\text{lk}_i = 1 \Rightarrow \text{false},$$

which implies that the procedure can return *true* or *false* when $\text{lk}_i \neq 1$. At first glance, this specification looks suspiciously weak, but a closer examination of the benchmark shows that lk_i is always

not equal to 1 when the checking code is invoked. Thus, our tool synthesizes a specification that returns any value when lk_i is 1.

Results Figure 4 shows the amount of time (in seconds) multi-abduction takes as we increase the number of locks. The original benchmark suite contains programs with 5 to 15 locks; here, we extend the suite with new programs with up to 60 locks. As with the parameterized benchmarks, only one multi-abduction call is made. The x -axis shows the size of the formula as we increase the number of locks. Within the short time limit of 3 minutes, our multi-abduction solver is able to scale to problems with ~ 5000 Boolean connectives, ~ 160 abducibles, and ~ 1500 quantified variables (generated from a program with 54 locks). Observe the need for multi-abduction here: We need to synthesize a specification for multiple procedures, one for each lock, and the same procedure might appear multiple times along control-flow paths emitted by the verifier. We would like to note that this set of benchmarks was specially crafted as a challenge to software verification tools: they require precise path reasoning; thus, a symbolic representation of paths is needed. Here, we do not only verify these benchmarks, we go one step further and synthesize maximal specifications for unknown code.

7.1.3 Summary

Our results demonstrate our multi-abduction procedure’s ability to scale to large formulas with thousands of Boolean connectives and hundreds of variables. Our experience with quantifier elimination indicates the importance of the order in which variables are eliminated. Going forward, we would like to investigate variable elimination ordering strategies, with an emphasis on our specification synthesis domain. Our tool is able to produce formulas in SMT-LIB format, which we will release to drive future research in decision procedures of quantified formulas and quantifier elimination.

7.2 E2: Specification synthesis for device drivers

Description We now turn our attention to our second set of experiments, which target device drivers. In device driver verification, stubs for OS kernel routines are typically manually written to avoid analysis of complex kernel code. Indeed, this is the case for the Static Driver Verifier (SDV) [39], Microsoft’s commercial product for driver developers, which uses the bounded verifier Corral [40].

In this section, we study the applicability of our proposed technique for automatically synthesizing specifications of kernel routines and initialization code for a standard suite of Windows drivers that are part of the software verification competition (SV-COMP).

		Multi-abduction statistics		
Driver/routine	Time	#Abds	DAG size	#Q-vars
kbfilter-s1				
init	1	7	498	113–120
IofCallDriver	0.1	1	137	21–31
IofComplRequest	0.03*			
KeSetEvent	0.1	1	104	21–26
KeWaitFSObject	0.1	1	123	19–32
kbfilter-s2				
init	3	7	987	236–243
IofCallDriver	0.1	1	165	29–39
IofComplRequest	0.04*			
KeSetEvent	0.1	1	132	29–34
KeWaitFSObject	0.1	1	151	27–40
floppy-s3				
init	5	6	1757	382–388
IofCallDriver	16	9	1398	241–317
IofComplRequest	8	4	1579	330–343
KeSetEvent	1	3	1551	318–333
KeWaitFSObject	2	4	1693	316–368
floppy-s4				
init	34	6	2767	628–634
IofCallDriver	27	11	2282	436–529
IofComplRequest	33	9	2535	548–576
KeSetEvent	3	3	2502	546–561
KeWaitFSObject	4	4	2644	544–596
diskperf				
init	145	7	2075	432–439
IofCallDriver	15	9	1356	220–303
IofComplRequest	15	5	1548	313–329
KeSetEvent	1	6	1573	305–335
KeWaitFSObject	5	6	1876	322–397
cdaudio				
init	184	6	5251	1172–1178
IofCallDriver	53	11	2972	596–696
IofComplRequest	175	17	3723	817–869
KeSetEvent	5	3	3526	795–810
KeWaitFSObject	19	8	3842	788–892

Table 1: Results for device drivers. Time is in seconds.

Here, we target the preprocessed benchmarks, from `ntdrivers-simplified` set, in which heap manipulation has been abstracted for the purposes of functional verification.

For evaluation, we identified all kernel routine stubs in these drivers: `IofCallDriver`, `IofCompleteRequest`, `KeWaitForSingleObject`, and `KeSetEvent`. These routines deal with I/O management and other core kernel issues, and are invoked by all drivers in the `ntdrivers` benchmarks.³ Also, all drivers contain an initialization stub, `init`, invoked at the start of execution, to set initial values for global variables, flags, etc. For both the kernel routine stubs and the initialization stubs, we applied our technique to synthesize a maximal specification that ensures driver correctness, with respect to properties encoded in the drivers. That is, we removed the stubs and attempted to automatically synthesize their specifications from the routines’ usage contexts, and we asked the following questions: (i) How efficient is the synthesis process? (ii) Do the synthesized specifications capture all behaviors of (i.e., over-approximate) the manually written stubs?

Synthesizing safe initialization specifications Synthesizing the initialization procedure of a device driver is an instance of input-filter synthesis: We infer the set of parameters for which the device driver behaves correctly. For each initialized global variable g in `init`, we replaced g ’s initialization code with $g \leftarrow \text{get}_g()$, where get_g is an unknown procedure. Thus, by synthesizing a maximal specification for each get_g , we know the range of values that g must

lie within in order to ensure correctness of the driver. Note that, for `init` specifications, we utilize multi-abduction to synthesize a specification per initialized variable. Simple abduction alone can only compute a monolithic formula over all uninitialized variables.

Table 1 shows the time it took to synthesize the specifications for all global variables for each driver in our set. The table also shows statistics for the final multi-abduction query: the number of abducibles (**#Abds**), the size of the problem (**DAG size**), and the number of quantified and total variables in the first call to `ABDUCE` (**#Q-vars**). For `init`, the number of abducibles (**#Abds**) corresponds to the number of initialized variables. For each driver, our synthesized specifications for procedures get_g were one of the following: (i) a specific value for g , equal to the one in the initialization routine; (ii) a range of values for g that includes the one in the initialization routine; or (iii) the specification `true`, which indicates that g does not affect the correctness of the procedure. For all the drivers, our techniques was able to compute a maximal specification for 6-7 procedures in a small amount of time: 1 to 184 seconds.

Synthesizing kernel routine stubs In each driver, we replaced each kernel routine stub with an unknown procedure whose arguments are the arguments of the kernel routine and the global variables read by the routine; the return values of the unknown procedure are its actual return value and values for all global variables it updates. Using our approach, we synthesized maximal safe specifications for these procedures.

Table 1 shows the amount of time it took to synthesize a specification for each kernel routine. (Routine `IofCompleteRequest` is not required for correctness of drivers `kbfilter-s*`, i.e., its maximal specification is `true`; this is detected by our front-end.) Consider the driver `cdaudio`; our technique computes a specification for `KeWaitForSingleObject` within 15 seconds, having to solve the (final) multi-abduction query containing ~ 3000 Boolean operators and ~ 800 quantified variables. We manually inspected synthesized specifications and in all cases found that they are equivalent to the original stubs or allow strictly more behavior. For instance, for the routine `IofCallDriver` in `kbfilter-s1`, our technique synthesized the disjunctive specification that states that `IofCallDriver` either updates the global variable `s` to the value of global variable `NP`, or returns the signal 259 and updates the value of `s` to that of `MPR3`. The specification does not mention the global variable `lowerDriverReturn`, which is updated in the given stub, meaning that its value is not needed for the correctness of the driver.

Summary Our evaluation indicates our technique’s ability to synthesize maximal safe specifications of unknown procedures in realistic programs. We have demonstrated this on a standard set of Windows driver benchmarks, where, within a short amount of time, we synthesized specifications of kernel routines that over-approximate existing stubs. We believe that our results demonstrate the promise of our approach, and provide strong evidence that it is indeed possible to synthesize specifications of unknown procedures.

8. Related Work

In this section, we place our specification synthesis technique in the context of related work.

Specification inference There is a large body of work on learning procedure specifications [7, 9, 12, 13, 20, 27, 35, 41, 45, 48, 50, 53, 54]. Many of these techniques [7, 9, 13, 45, 48, 50, 53] use either client or library code to learn API-usage rules of libraries, which can then be used for verifying other clients of the same library. Here, we addressed an orthogonal problem: synthesizing logical specifications of unknown procedures that ensure client safety.

The works of Bastani et al. [12] and Zhu et al. [54] have a similar goal to ours, but in the context of taint-flow analysis. Bastani et al.

³ See the Windows Driver Kit (WDK) for the detailed API [3].

present a CFL-based solution for computing taint-flow and aliasing specifications. Similarly, Zhu et al. use simple abduction to find the smallest number of assumptions on unknown procedures. In contrast to these works, our technique differs in two ways. First, we present a general logic-based technique that is not tied to taint-flow analysis or CFL-based analyses. Second, our technique ensures a universally maximal solution, whereas maximality of prior work is relative to an initial taint-flow analysis, which is likely imprecise.

Seghir and Kroening [49] and Cimatti et al. [23] compute a weakest safe precondition of a program or transition system. This problem can be modelled in our framework by inserting a single unknown procedure call at the beginning of the program, where the procedure call returns initial values for program parameters.

Work on *angelic verification* [20, 27] aims to find *reasonable* environment assumptions in order to suppress “*stupid*” false alarms. In the recent paper by Das et al. [27], the user is assumed to specify a set of acceptable specifications; the job of the verifier is to find a specification within the supplied set that suppresses bugs while being permissive (limits dead code). Our work is similar in that we look for specifications of unknown code that make the program correct. However, our work differs in that we present a general logical treatment of the problem. In particular, our multi-abduction algorithm enables automatic computation of maximal specifications, but *without user intervention*.

Abduction in analysis and verification Abduction made its way into computer science through artificial intelligence, where it was used for knowledge representation and reasoning [29, 43]. Here, we focus on use of abduction in program analysis and verification.

In the context of separation logic for shape analysis, *bi-abduction* was introduced in the seminal work of Calcagno et al. [22] as a means for inferring procedure preconditions that ensure memory-safe execution. Our work differs in a number of ways: First, we target first-order theories that admit quantifier elimination. Second, multi-abduction finds specifications of unknown procedures in addition to preconditions, whereas bi-abduction assumes that it has access to callee specifications. We believe that our approach can be profitably combined with bi-abduction in a bottom-up compositional reasoning fashion. For instance, the recently open-sourced tool, Infer [1], a commercial version of Calcagno et al.’s technique, does not perform precise reasoning over arithmetic operations, and may return false positives in such cases. By incorporating a first-order abduction mechanism alongside shape abduction, we can create a more precise analysis technique.

The work of Qin et al. [47] presents a combined shape-data abstract domain that can be used for inferring specifications of unknown procedures using a form of abduction over the abstract domain. In comparison, we (i) compute maximal specifications of unknown procedures, whereas Qin et al.’s work heuristically abduces a specification; and (ii) our abduction queries always find a solution if one exists, as we are not restricted by an abstract domain that loses information during forward propagation.

In program analysis, abduction appeared first in the work of Giacobazzi [31] in the context of logic programming. Giacobazzi defined a top-down abstract interpretation of logic programs that abduces properties of modules that ensure correctness of their composition. Our work is analogous, in the sense that we also ask, “*How should callees behave in order to ensure correctness?*” In comparison to Giacobazzi’s work, we target imperative programs and do not constrain our analysis to an abstract domain; instead, we encode program semantics as first-order formulas and utilize multi-abduction to synthesize weakest possible procedure specifications necessary for ensuring correctness.

Horn-clause solving Our work shares similarities with works on solving *Horn clauses* [15, 32, 33]. Our technique can be viewed as

a solver for *recursive* Horn clauses with *unconstrained* predicates: those that appear exclusively in the bodies of clauses, and not as heads. Intuitively, unconstrained predicates represent specifications of unknown procedures—computed by multi-abduction—and constrained predicates represent inductive invariants proving that the program is safe—computed by the verifier. Existing work on solving Horn clauses does not address the question of optimality and would allow trivial solutions that set unconstrained predicates to *false*, as allowed by Horn-clause semantics. In contrast, our multi-abduction solver ensures that we compute maximal, non-trivial solutions. Picking up a recursive Horn-clause solver—like HSF [32] or GPDR [36]—and applying it in our setting does not guarantee a useful solution. Indeed, in our experience, unconstrained predicates in GPDR are always set to *false*. Thus, our contribution in the context of Horn-clause solving is a recursive Horn-clause solver that ensures non-trivial, maximal solutions to unconstrained predicates.

Beyene et al. [14] present a game solving technique that is formalized as Horn-clause solving. The authors demonstrate their approach for program repair and synthesis of winning strategies in two-player games. Our approaches are closely related. For instance, in Beyene et al., program repair is formalized as finding a strengthening of a program’s transition relation. In our setting, this is equivalent to finding a specification of an unknown procedure that restricts program executions to safe ones. In comparison with Beyene et al., our work does not require user-supplied linear-arithmetic templates for specifications of unknown procedures; instead, it synthesizes a maximal specification over a first-order theory, a much more challenging problem. On the other hand, Beyene et al. also handle temporal specifications, allowing them to stipulate that a synthesized procedure specification ensures total correctness.

Formula decomposition Veanes et al. introduced *monadic decomposition* [52], a technique that takes a formula $F(x, y)$ and computes k pairs of formulas $(A_i(x), B_i(y))$ such that $\bigvee_i (A_i(x) \wedge B_i(y)) \equiv F(x, y)$. Our formulation of Cartesian decomposition is similar to 1-monadic decomposition, but it is more general: We relax the equivalence constraint to $A(x) \wedge B(y) \Rightarrow \varphi$ and look for maximal solutions of A and B . Also, Cartesian decomposition does not restrict the decomposition to a pair of formulas, but allows an arbitrary number of unknowns. In addition, we consider isomorphic decomposition. It is also important to note that our Cartesian decomposition procedure computes a 1-monadic decomposition when it exists.

9. Conclusion

We studied the problem of synthesizing maximal safe specifications of unknown procedures and presented a logic-based solution that leverages multi-abduction. Our empirical evaluation demonstrates the effectiveness of this technique at finding useful procedure specifications for realistic programs.

We believe there are many exciting and important applications of our work; in the future, we plan to investigate the following:

- **Bottom-up compositional analysis for shape-data:** As discussed in Section 8, our abduction algorithms can be combined with the bi-abduction-based shape analysis of Calcagno et al. [22] to produce a combined shape-data analysis. In the future, we plan to study such a combination with the goal of reducing false positives in shape analysis.
- **Compositional top-down verification:** We also plan to investigate the effectiveness of our technique for top-down verification. We can use our technique to verify callers independently of callees and then separately check whether the callees satisfy the inferred maximal specifications. In comparison with forward abstract-interpretation-based, bottom-up, and

interpolation-based analyses, this direction has received very little attention [31, 47].

Acknowledgements We would like to thank Kenneth McMillan for one of the examples, and Zachary Kincaid, Thomas Reps, and Ben Liblit for comments on an early draft of the paper. The second author is funded by NSF Award #1453386 and FA 8750-14-2-0270.

References

- [1] Infer. <http://fbinfer.com/>.
- [2] The LLVM compiler infrastructure. <http://llvm.org>.
- [3] Windows driver kit (WDK). [https://msdn.microsoft.com/en-us/library/windows/hardware/ff557573\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/hardware/ff557573(v=vs.85).aspx).
- [4] A. Albarghouthi, A. Gurfinkel, and M. Chechik. Craig interpretation. In *SAS*, 2012.
- [5] A. Albarghouthi, A. Gurfinkel, and M. Chechik. From Under-approximations to Over-approximations and Back. In *TACAS*, 2012.
- [6] A. Albarghouthi, Y. Li, A. Gurfinkel, and M. Chechik. UFO: A framework for abstraction-and interpolation-based software verification. In *CAV*, 2012.
- [7] R. Alur, P. Černý, P. Madhusudan, and W. Nam. Synthesis of interface specifications for Java classes. In *POPL*, 2005.
- [8] R. Alur, R. Bodik, G. Juniwal, M. M. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *FMCAD*, 2013.
- [9] G. Ammons, R. Bodík, and J. R. Larus. Mining specifications. In *POPL*, 2002.
- [10] T. Ball and S. K. Rajamani. The SLAM toolkit. In *CAV*, 2001.
- [11] M. Barnett and K. R. M. Leino. Weakest-precondition of unstructured programs. In *PASTE*, 2005.
- [12] O. Bastani, S. Anand, and A. Aiken. Specification inference using context-free language reachability. In *POPL*, 2015.
- [13] N. E. Beckman and A. V. Nori. Probabilistic, modular and scalable inference of typestate specifications. In *PLDI*, 2011.
- [14] T. Beyene, S. Chaudhuri, C. Popeea, and A. Rybalchenko. A constraint-based approach to solving games on infinite graphs. In *POPL*, 2014.
- [15] T. A. Beyene, C. Popeea, and A. Rybalchenko. Solving existentially quantified horn clauses. In *CAV*, 2013.
- [16] D. Beyer. Status report on software verification - (Competition summary SV-COMP 2014). In *TACAS*, 2014.
- [17] D. Beyer, T. A. Henzinger, R. Jhala, and R. Majumdar. The software model checker BLAST. *STTT*, (5-6), 2007.
- [18] D. Beyer, A. Cimatti, A. Griggio, M. E. Keremoglu, and R. Sebastiani. Software Model Checking via Large-Block Encoding. In *FMCAD*, 2009.
- [19] N. Bjørner. Linear quantifier elimination as an abstract decision procedure. In *IJCAR*, 2010.
- [20] S. Blackshear and S. K. Lahiri. Almost-correct specifications: A modular semantic framework for assigning confidence to warnings. In *PLDI*, 2013.
- [21] D. Brumley, H. Wang, S. Jha, and D. X. Song. Creating vulnerability signatures using weakest preconditions. In *CSF*, 2007.
- [22] C. Calcagno, D. Distefano, P. O'Hearn, and H. Yang. Compositional shape analysis by means of bi-abduction. *POPL*, (1), 2009.
- [23] A. Cimatti, A. Griggio, S. Mover, and S. Tonetta. Parameter synthesis with IC3. In *FMCAD*, 2013.
- [24] E. Clarke, D. Kroening, and F. Lerda. A Tool for Checking ANSI-C Programs. In *TACAS*, pages 168–176, 2004.
- [25] D. Cooper. Theorem proving in arithmetic without multiplication. *Machine Intelligence*, (91-99), 1972.
- [26] M. Costa, M. Castro, L. Zhou, L. Zhang, and M. Peinado. Bouncer: Securing software by blocking bad input. In *SOSP*, 2007.
- [27] A. Das, S. K. Lahiri, A. Lal, , and Y. Li. Angelic verification: Precise verification modulo unknowns. In *CAV*, 2015.
- [28] L. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, 2008.
- [29] M. Denecker and A. C. Kakas. Abduction in logic programming. In *Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part I*, 2002.
- [30] I. Dillig, T. Dillig, and A. Aiken. Automated error diagnosis using abductive inference. *PLDI*, 2012.
- [31] R. Giacobazzi. Abductive analysis of modular logic programs. In *ISLP*, 1994.
- [32] S. Grebenschikov, A. Gupta, N. P. Lopes, C. Popeea, and A. Rybalchenko. HSF (C): a software verifier based on Horn clauses. In *TACAS*, 2012.
- [33] S. Grebenschikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In *PLDI*, 2012.
- [34] S. Gulwani. Synthesis from examples. *WAMBSE*, (2), 2012.
- [35] T. A. Henzinger, R. Jhala, and R. Majumdar. Permissive interfaces. In *ESEC/SIGSOFT FSE*, 2005.
- [36] K. Hoder and N. Bjørner. Generalized property directed reachability. In *SAT*, 2012.
- [37] B. Jobstmann, A. Griesmayer, and R. Bloem. Program repair as a game. In *CAV*, 2005.
- [38] V. Kuncak, M. Mayer, R. Piskac, and P. Suter. Complete functional synthesis. In *PLDI*, 2010.
- [39] A. Lal. personal communication, 2015.
- [40] A. Lal, S. Qadeer, and S. Lahiri. Corral: A solver for reachability modulo theories. In *CAV*, 2012.
- [41] B. Livshits, A. V. Nori, S. K. Rajamani, and A. Banerjee. Merlin: Specification inference for explicit information flow problems. In *PLDI*, 2009.
- [42] F. Long, S. Sidiroglou-Douskos, D. Kim, and M. Rinard. Sound input filter generation for integer overflow errors. 2014.
- [43] S. McIlraith. Logic-based abductive inference. Technical Report KSL-98-19, Knowledge Systems Laboratory, July 1998.
- [44] K. L. McMillan. Lazy abstraction with interpolants. In *CAV*, 2006.
- [45] J. W. Nimmer and M. D. Ernst. Automatic generation of program specifications. In *ISSTA*, 2002.
- [46] W. Pugh. The Omega Test: a fast and practical integer programming algorithm for dependence analysis. *CACM*, 1992.
- [47] S. Qin, C. Luo, G. He, F. Craciun, and W. Chin. Verifying heap-manipulating programs with unknown procedure calls. In *ICFEM*, 2010.
- [48] M. K. Ramanathan, A. Grama, and S. Jagannathan. Static specification inference using predicate mining. *PLDI '07*, 2007.
- [49] M. N. Seghir and D. Kroening. Counterexample-guided precondition inference. In *ESOP*, 2013.
- [50] S. Shoham, E. Yahav, S. Fink, and M. Pistoia. Static specification mining using automata-based abstractions. *ISSTA*, 2007.
- [51] A. Solar-Lezama. *Program Synthesis By Sketching*. PhD thesis, EECS Department, University of California, Berkeley, Dec 2008.
- [52] M. Veanes, N. Bjørner, L. Nachmanson, and S. Bereg. Monadic decomposition. In *CAV*, 2014.
- [53] J. Yang, D. Evans, D. Bhardwaj, T. Bhat, and M. Das. Perracotta: mining temporal api rules from imperfect traces. *ICSE*, 2006.
- [54] H. Zhu, T. Dillig, and I. Dillig. Automated inference of library specifications for source-sink property verification. In *APLAS*, 2013.

A. Proofs

Proof of Theorem 1

Direction $\Rightarrow$ Let C be an abducible-free formula and let φ be a formula containing abducibles only positively. (We assume that C 's free variables do not include those fresh variables introduced by FLATTEN.) Let Δ be a non-trivial mapping from each abducible in φ to a formula such that $\varphi[\Delta] \Rightarrow C$ is valid. Further, suppose that FLATTEN(φ) yields $\psi \wedge \varphi^*$, where ψ is a conjunction of abducibles. Consider any model M of the formula $\varphi[\Delta] \Rightarrow C$. Since $M \models \varphi[\Delta] \Rightarrow C$, we have either

- (i) $M \not\models \varphi[\Delta]$ or
- (ii) $M \models \varphi[\Delta]$ and $M \models C$.

Suppose (i). Observe that $\varphi[\Delta] \equiv \exists X_f. \psi[\Delta] \wedge \varphi^*$, where X_f is the set of fresh variables introduced by FLATTEN. Hence, if $M \not\models \varphi[\Delta]$, then no extension M' of M can entail $\psi[\Delta] \wedge \varphi^*$.

For (ii), if $M \models \varphi[\Delta]$, then let M' be an extension of M that assigns c to x_i whenever M assigns c to x (where x_i are the fresh variables corresponding to the variables x). We then have $M' \models (\psi[\Delta] \wedge \varphi^*)$ and $M' \models C$. Hence, $M' \models (\psi[\Delta] \wedge \varphi^*) \Rightarrow C$. For any other extension M'' of M' , we have $M'' \not\models \varphi^*$; hence $M'' \models (\psi[\Delta] \wedge \varphi^*) \Rightarrow C$.

Direction $\Leftarrow$ We omit this direction of the proof, as it is similar to the first part.

Proof of Theorem 2

Non-triviality The fact the returned specification is non-trivial follows from the definition of multi-abduction (condition (ii) of Definition 5).

Safety Safety of the specification simply follows from the check at lines 5-6, which ensures that $P[\Delta]$ satisfies $\text{post}(P)$.

Maximality The maximality of Δ follows from Theorem 1 and the maximality of the solution computed by MULTIABDUCE (Theorems 3 and 4). In particular, suppose there was a better solution Δ' such that $\Delta' \succ \Delta$ and $P[\Delta'] \models \text{post}(P)$. For Δ' to be safe, it would need to rule out all executions in Π in the last iteration of the algorithm; that is, it would need to be a solution to the multi-abduction problem defined by Φ and $\text{post}(P)$. However, this means that $\Delta' \not\models \Delta$; otherwise, it would violate the maximality of the solution computed by MULTIABDUCE.

Proof of Theorem 3

Case 1: solvable problem If the problem has a solution, then the ABDUCE call in MULTIABDUCEUNIQUE returns an *upper-bound* ψ such that for any solution Δ , $\bigwedge_i \Delta(R_i) \Rightarrow \psi$ is valid. Note also that, since ABDUCE succeeds, by definition of ABDUCE, we know that CARTECOMP will find an initial solution when calling INIT SOLN.

The fact that CARTECOMP weakens one $\Delta(R_i)$ at a time, starting from an initial solution, guarantees that the result is maximal. Suppose it was possible to weaken the solution for one of the R_i 's while fixing the rest to $\Delta(R_j)$. Clearly, the weakest solution for R_i cannot be weaker than $\text{ABDUCE}(\bigwedge_{i \neq j} \Delta(R_j), \psi, x_i)$. However, since the algorithm iteratively weakens every R_j , the solution $\Delta(R_j)$ at the end is weaker than (or at least as weak as) the solution used when computing R_i . Hence, R_i cannot be weakened beyond $\Delta(R_i)$.

Case 2: unsolvable problem Suppose the problem has no solution, then the ABDUCE call in MULTIABDUCEUNIQUE fails and MULTIABDUCEUNIQUE returns none. By definition of simple abduction, there is no solution of our multi-abduction problem with

unique unknowns. If simple abduction succeeds, then INIT SOLUTION must succeed and by the argument above MULTIABDUCEUNIQUE finds a maximal solution.

Proof of Theorem 4

Case 1: returns solution (Sketch) It is easy to see that the mapping Δ returned by MULTIABDUCE is a solution. To see why it is maximal, suppose it was possible to weaken some R_i while fixing the other R_k 's. This implies that there must be some φ_i over x_i such that $\Delta(R_i) \Rightarrow \varphi_i$ and

$$\left(\bigwedge_j \varphi_i[x_{ij}/x_i] \wedge \bigwedge_{k \neq i} \bigwedge_j \Delta(R_k)[x_{kj}/x_k] \right) \Rightarrow \psi$$

Since ABDUCE yields a weakest solution, we must have:

$$\bigwedge_j \varphi_i[x_{ij}/x_i] \Rightarrow \text{ABDUCE} \left(\bigwedge_{k \neq i} \bigwedge_j \Delta(R_k)[x_{kj}/x_k], \psi, x_i \right)$$

Let's call the result of the ABDUCE call above γ . Now, consider the solution φ_k for each R_k at the time we computed the solution for unknown R_i at lines 8-9. Since the algorithm only weakens solutions, we have $\varphi_k \Rightarrow \Delta(R_k)$ which implies $\gamma \Rightarrow \psi_i$ where ψ_i is the value at line 8 when we computed the solution for R_i . Hence, we have $\bigwedge_j \varphi_i[x_{ij}/x_i] \Rightarrow \psi_i$. Since φ_i is weaker than $\Delta(R_i)$, this in turn implies $\text{SYMDECOMP}(\psi_i, \mathcal{A}(R_i), \Delta') \Rightarrow \varphi_i$ where Δ' is the solution at the time we found an interpretation for R_i . However, this would mean that $\Delta(R_i)$ is not a solution to the SYMDECOMP problem because it is possible to find a new model of the formula at line 21.

Case 2: returns no solution If MULTIABDUCE returns none, then it is because the call to simple ABDUCE in MULTIABDUCE did not find a solution. By definition of simple abduction, this means that there is no solution the multi-abduction problem.

Proof of Theorem 5

Consider the following multi-abduction problem in linear real arithmetic (LRA): $R(x) \wedge R(y) \Rightarrow y \neq x + 1$. We prove that this problem has no maximal solution.

First, observe that for any (maximal or non-maximal) solution $[R(\alpha) \mapsto \varphi]$, it must be the case that for any $c, k \in \mathbb{R}$, where $k > 1$, the implication $c < \alpha \leq c + k \Rightarrow \varphi$ is not logically valid. Intuitively, this states that there is no contiguous region of models of φ that is an interval of size ≥ 1 , as stipulated by the right-hand side of the multi-abduction problem. Since φ is an LRA formula over the single real variable α , it represents a finite set of polyhedra in $\mathbb{R}^1$ where each polyhedron is of size < 1 .

Using this fact, we now show that, for any solution φ of the problem, there is a strictly weaker solution φ' , thus showing that there is no maximal solution. Let us fix a solution φ . Now, there must be constants $c, k \in \mathbb{R}$, where $0 < k \leq 1$, such that (i) $c < \alpha < c + k \vee \varphi$ is also a solution to the problem, and (ii) $c < \alpha < c + k \not\Rightarrow \varphi$. Suppose this is not true, then either φ contains infinitely-many polyhedra, which cannot be represented in LRA, or there is an open polyhedron in φ , which cannot be case since polyhedra are restricted to intervals of size < 1 , and therefore must be bounded.

It follows, by contradiction, that there are problems in LRA for which there are no maximal solutions.

B. Synthesizing for One Path at a Time

In Algorithm 1, in each iteration we synthesize a specification that works for all paths in Π , and Π grows monotonically, so the multi-abduction queries grow in size.

A natural question to ask is whether we can refine the current specification one path at a time. For instance, can we simply use the specification computed for the most recent path to strengthen the current specification? To see why this would not work, consider the following program (we encode the postcondition as assertions, and use $*$ to denote non-deterministic choice).

```
x ← foo()
y ← bar()
if (*)
  assert(x + y ≥ 0)
else
  assert(x + y ≤ 0)
```

Suppose VERIFY returns the path that goes through the then branch in the first iteration. Multi-abduction might return the maximal solution that maps `foo` to $\beta_f \geq 10$ and `bar` to $\beta_b \geq -10$, where β_f and β_b are the return values of `foo` and `bar`, respectively.

Now, in the next iteration, VERIFY will return the path that goes through the else branch. If we consider this path in isolation of the previous path, multi-abduction might compute the maximal solution $\beta_f \leq 0$ and $\beta_y \leq 0$. This contradicts our previously computed specification. Thus, to compute a maximal specification, we need to consider all paths seen so far.

C. Dealing with Failed Multi-abduction Calls

In Algorithm 1, multi-abduction might fail to find a solution for the current set of paths Π . If this is the case (line 11), this means that the only specifications that make P_Π safe are those that make all paths in Π infeasible—by setting some specifications to *false* or contradicting path conditions. Of course, we could use such specifications and continue with our counterexample-guided loop. But this may lead to specifications that make every single path through the program infeasible. Consider the following example for an illustration:

```
x ← foo()
y ← bar()
assume (x < 0)
if (y > 0)
  err ← true
else
  err ← false
```

Suppose the postcondition is $\neg err$ and that we want to find specifications for `foo` and `bar`. Starting with the specification *true* for both procedures, VERIFY first finds the counterexample that goes through the then branch of the conditional. Any safe specification for `foo` and `bar` must contradict the condition $x < 0$ or the condition $y > 0$. As a result, multi-abduction will fail, since any solution will be inconsistent with the left-hand side of the abduction problem (Φ from Algorithm 1). But suppose we relax this condition of multi-abduction and allow solutions that contradict the left-hand side. In this example, we might compute the specification $\beta_f > 0$ for `foo` and *true* for `bar` (where β_f denotes the return value of `foo`). This is a safe and maximal specification for the path under consideration; it is also a safe maximal specification for the whole program. Unfortunately, it makes both paths through the program infeasible, as it contradicts the *assume* condition. What we want is the specification *true* for `foo` and $\beta_b < 0$ for `bar`, but since the only path we considered was the one that had to be infeasible, we missed this specification.

To avoid such scenarios, we make sure that some paths in the set Π have safe maximal specifications that do not make them infeasible. To do so, we make use of the set $\hat{\Pi}$, which stores paths that have to be infeasible and avoids finding them again in the verification phase.